

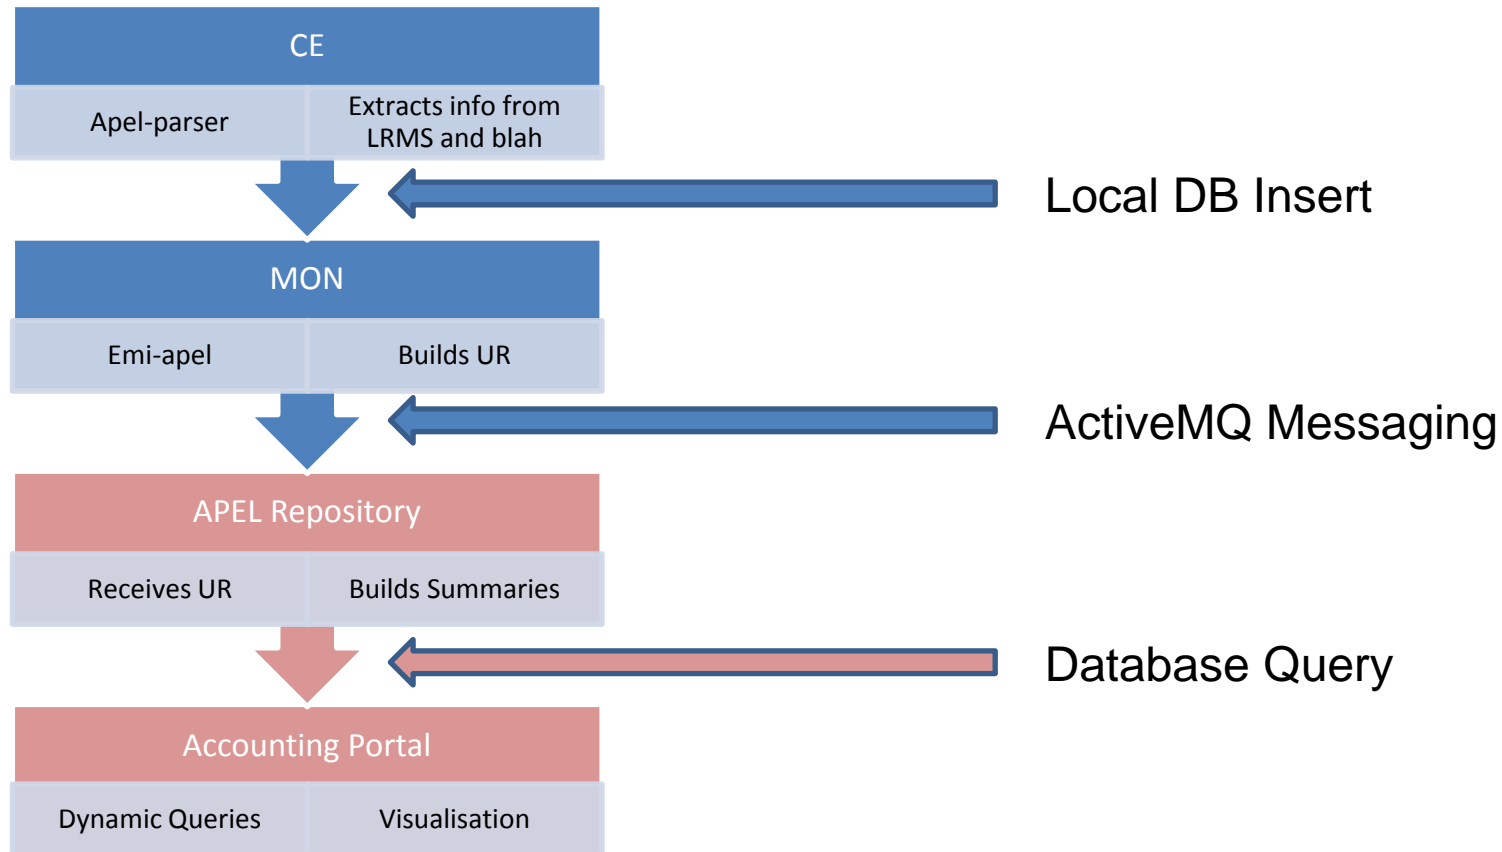
# Accounting Developments

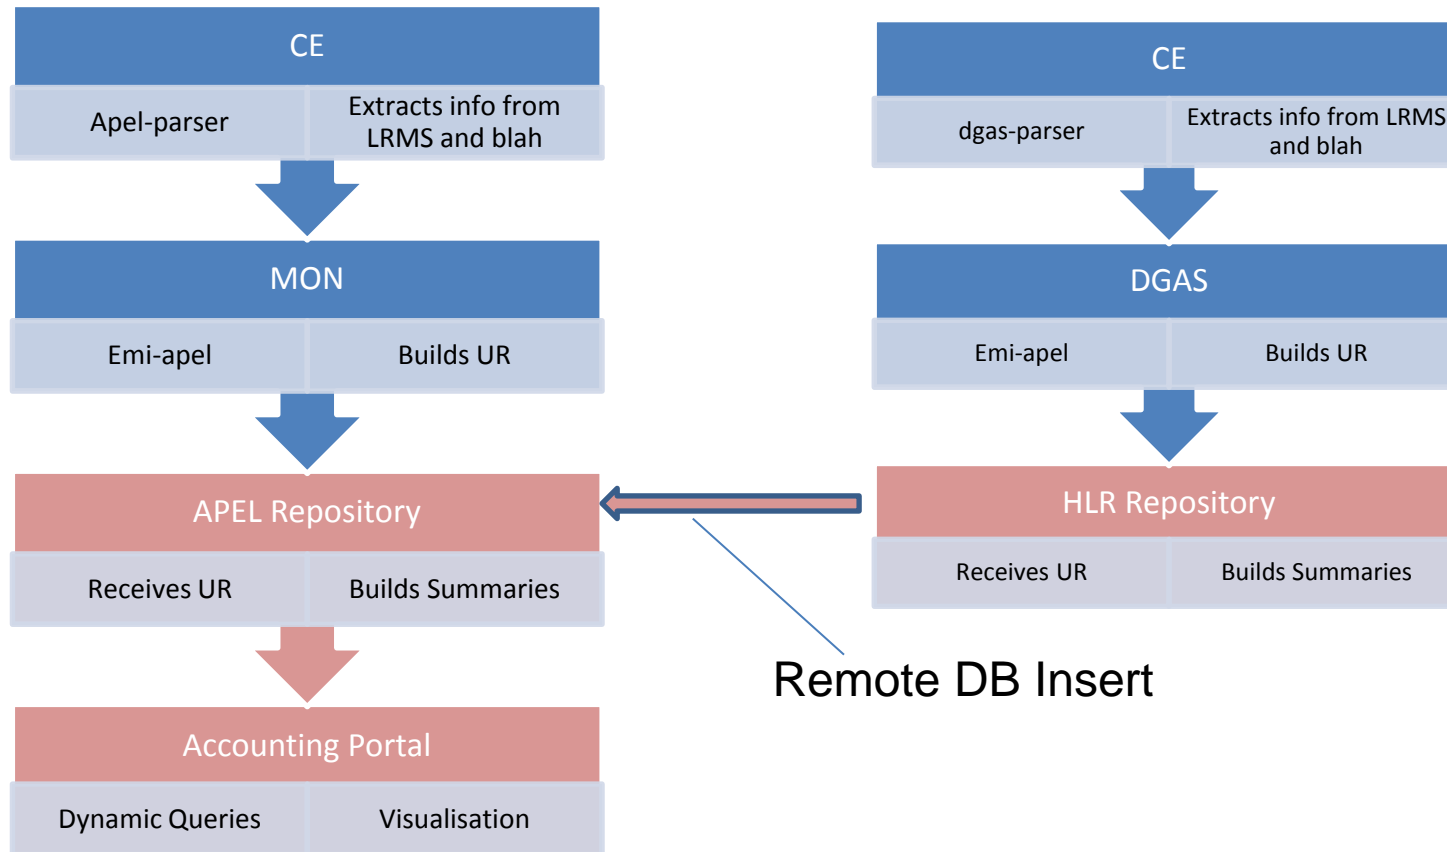
*From EMI Accounting standards to  
EGI Accounting Services  
and how they affect WLCG*

John Gordon

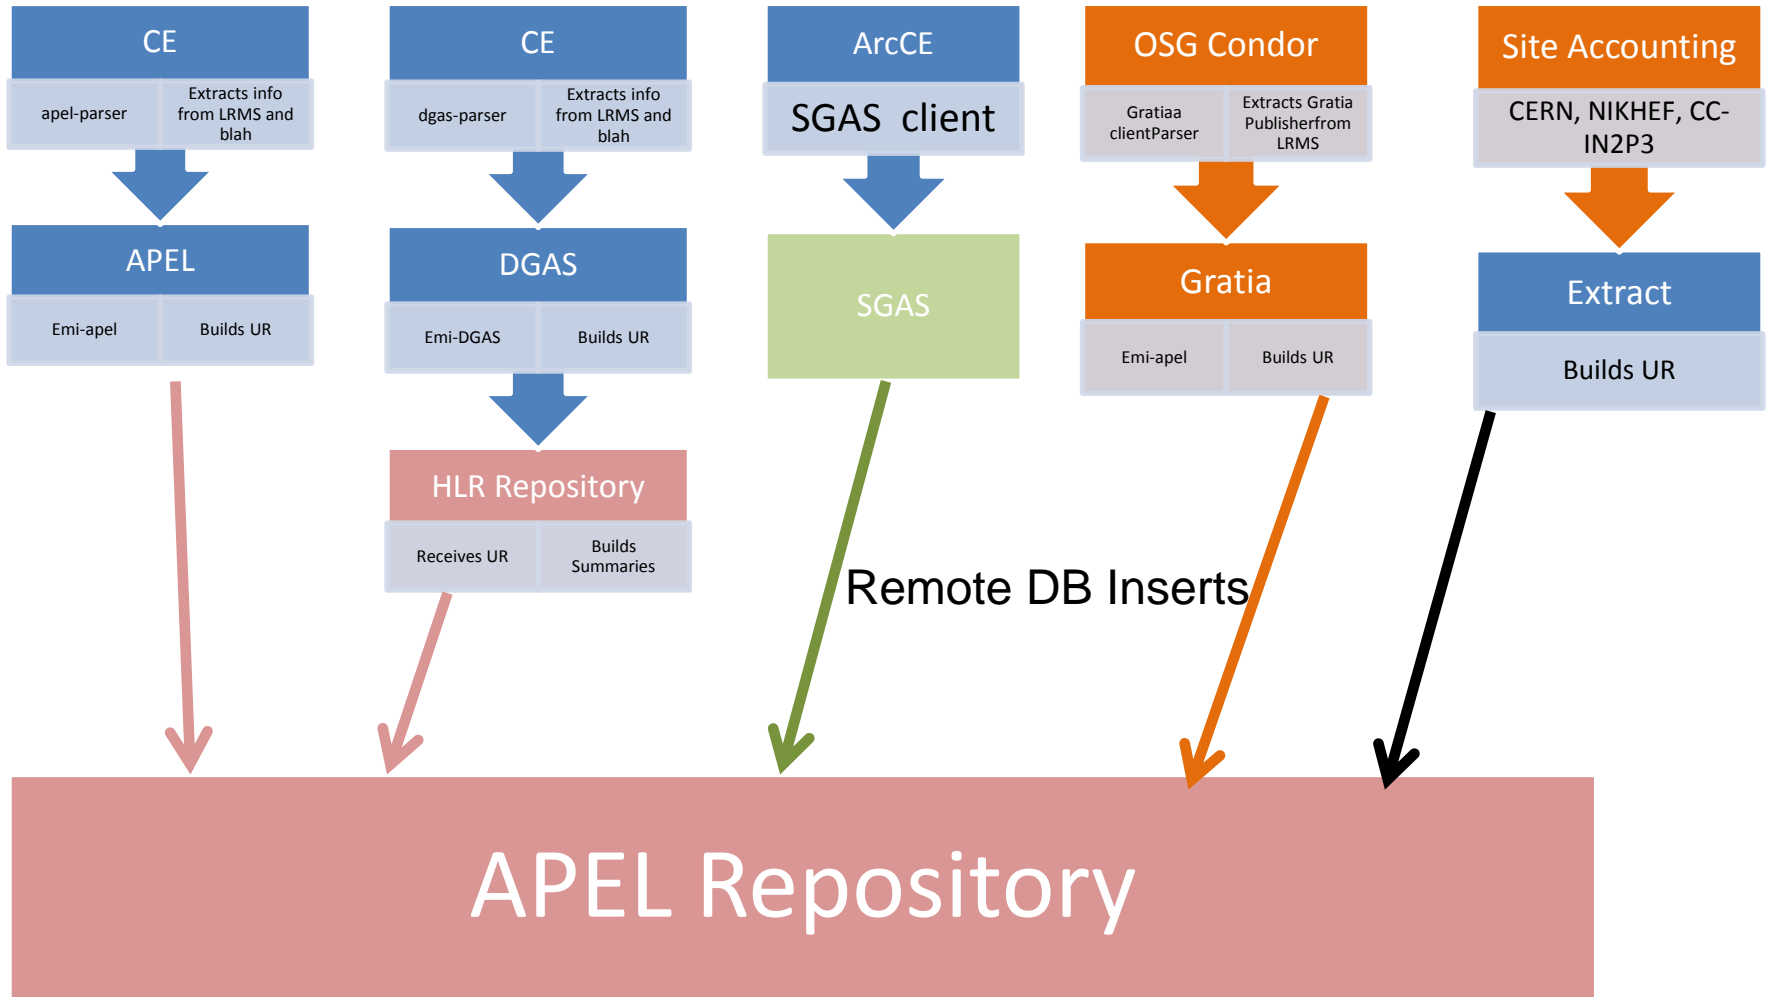
GDB, November 2011

- Existing EGI Infrastructure
- New version under development
- Extension to cover more types of site
- New Types of accounting





# Existing EGI Infrastructure




- EGI gathers accounting data in one place for:-
- All EGI VOs
- LHC Experiment VOs
- Other partner VOs (Australia, Asia-Pacific)

- Version 2 Messaging = SSM
  - Secure Stomp Manager

# What is the SSM?

- Transfers files between computers using Python, ActiveMQ and STOMP
- One SSM is needed at each end
  - Producer and Consumer



- If your SSMs are working, the files get sent reliably

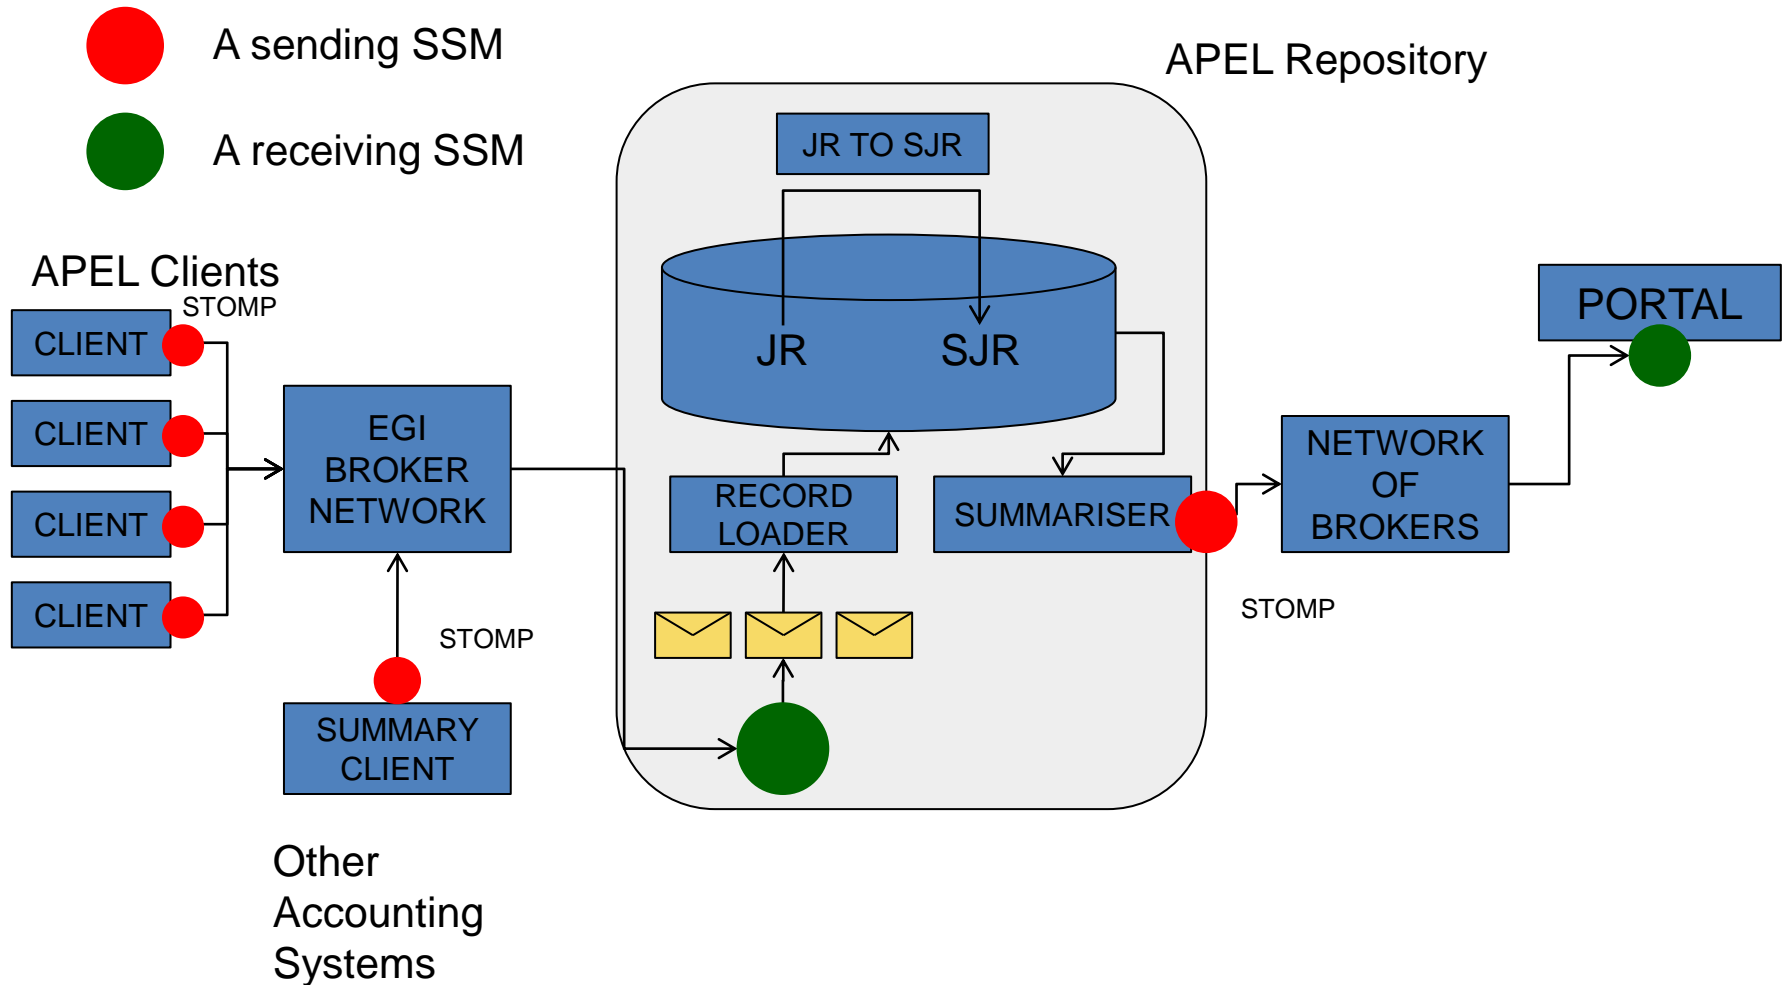


# Why is it good?

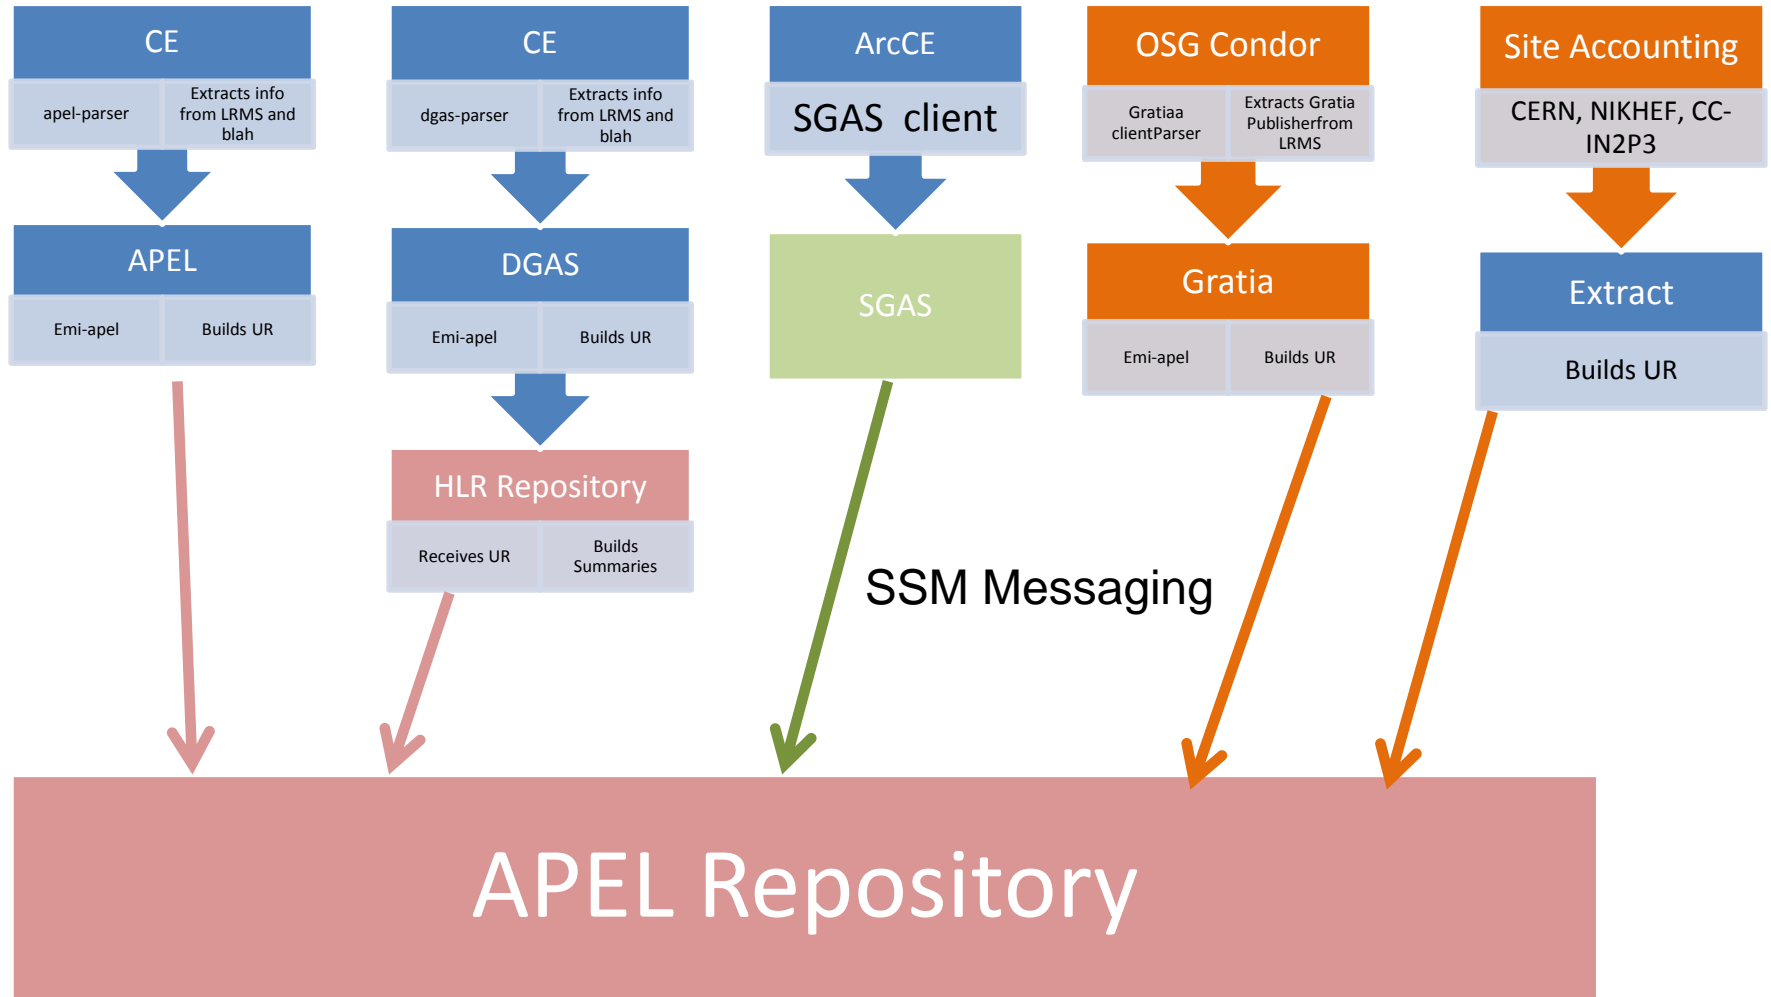
- Messaging 100% separated from storage
  - Re-usable
  - Easy to diagnose problems
- If the database is down – consumer still up
  - Incoming messages still written to disk
- Using EGI network of brokers provided
- STOMP protocol used by other tools – support available
- Using host certificates to sign and encrypt messages
- Messages are encrypted whenever they're between SSMs
  - EGI policy requirement for User data on the WAN

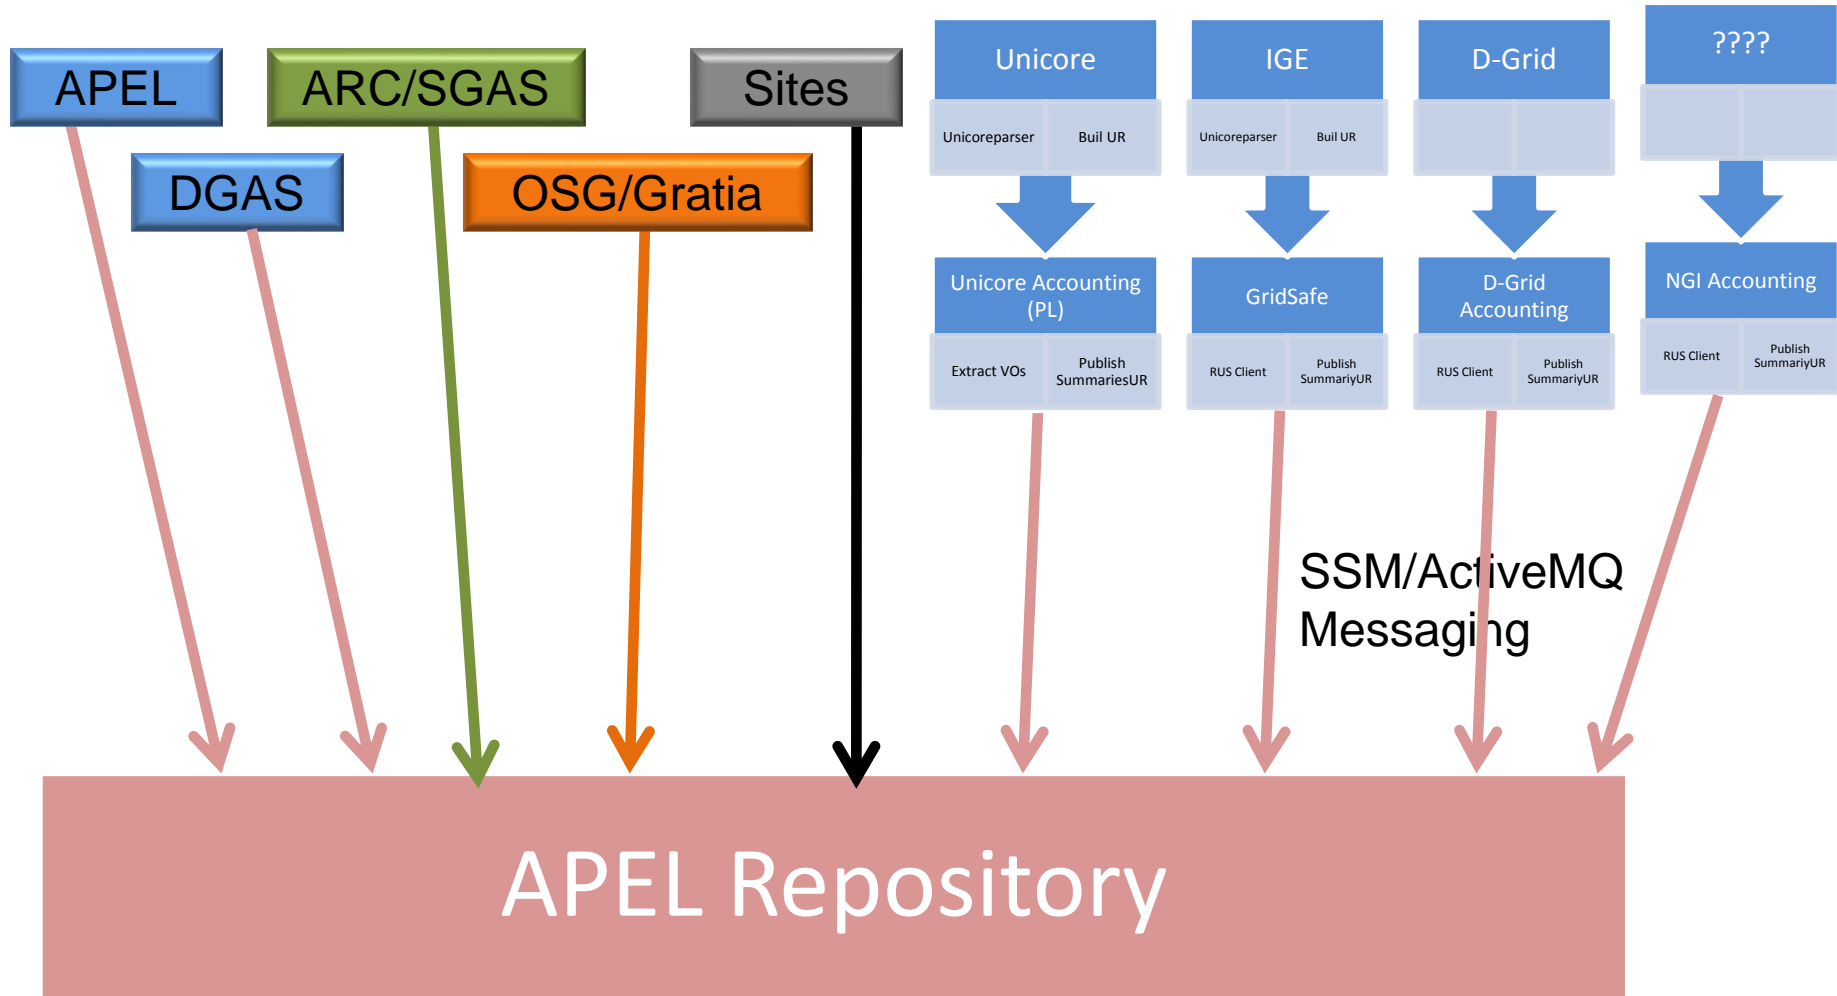
# Why is it good? (cont.)

- Producer and/or consumer mode (same program)
- Message transmission handling: retries, acknowledgements and ack-tracking
  - You know if a message has got through
  - A lost message is never a problem
- Access Control List of valid message producers for authentication
- Independent of message content
- Available To Others



- Existing systems (DGAS, SGAS, OSG, etc) will move to using SSM
  - Or the same format
  - Under test
- Distribute as a national accounting service
- Extend to Other Clients
  - Some under test already
  - CERN talk next





- The new EGI Infrastructure will leave us in a good position to extend it to accept other types of accounting record.
  - Configure consumer to vary workflow according to record type
  - Designed in to allow schema evolution
  - Different instances of consumer configured to handle different record types
  - Different queues/topics in messaging
- Propose SSM for new record types too.

- StAR – Storage Accounting Record
  - From EMI Data Group to OGF
  - Not yet compatible with the WLCG Installed Capacity use case.
- EMI Storage Services will gather data to populate StAR (dCache, dpm, StoRM)
- What then?
  - Work with them to publish UR into the infrastructure
- Evangelise to other storage products (Castor, EOS, BestMan)
- EMI aiming for 2013 release but should be widespread testing before then.



- MPI users pushing for consistent accounting of parallel jobs
- This is relevant to whole-node use and how one treats cpu/wallclock.
- Record nodecount and ncpus(cores) for batch jobs

EGI has ambitions to address accounting of:

- Clouds/Virtualisation
- Applications
- Data Centre Use
- MPI
- Services

Does WLCG Require anything else?

- Networking
- Memory
- Pilot Jobs?

- Decide what you want to account
- Define a draft Usage Record
- Get agreement
- Gather data to populate UR
  - For each service of relevant type
- Publish UR into an Accounting Infrastructure
  - ditto
- Decide and implement workflows for data
  - In receiving infrastructure
- Decide and implement reporting and visualisation

- New emi-APEL Client next year in EMI-2
  - Publish into new APEL infrastructure
  - No big change if you have already moved to emi-apel from glite-apel.
- LRMS (batch) Support
  - Any demand for Condor?
  - Historically weak for (S)GE
  - APEL Team support parsers for relevant batch systems but they need help from LRMS experts in parsing and interpreting batch logs

- Like all the other middleware which interfaces with batch systems, the APEL Parsers are difficult to configure automatically and sensitive to change.
- Requires manual configuration by site and maintenance of this
- Log locations. structure and lifetime
- Benchmarking
  - glite-cluster may help
- Other suggestions welcome

- Hierarchical Tree
- Tier 1
    - Tier 2
      - Australia
      - Austria
      - Belgium
      - Brazil
      - Canada
      - China
      - Czech Republic
      - Estonia
      - Finland
      - France
      - Germany
      - Greece
      - Hungary
      - India
      - Israel
      - Italy
      - Japan
      - Norway
      - Pakistan
      - Poland
      - Portugal
      - Republic of Korea
      - Romania
      - Russian Federation
      - Slovenia
      - Spain
      - Sweden
      - Switzerland
      - Taipei

EGI View -> Tier2

 Print Page

Data to graph:	Norm. Sum CPU (kSI2K-hours)	Normalised CPU time to a reference value of 1000 SpecInt2000		
Period:	Start year: 2011	Start month: 1	End year: 2011	End month: 11
Groupings:	Show data for: COUNTRY_T2	as a function of: VO		
VO Groups:	<input checked="" type="radio"/> LHC <input type="radio"/> TOP 10 <input type="radio"/> ALL <input type="radio"/> Custom <input checked="" type="checkbox"/> Group the rest of VOs in a new category			
Chart:	Type: ACCUM BAR	Scale: LINEAR		
dteam VO:	<input type="checkbox"/> Exclude dteam and ops VOs jobs information			

Refresh

COUNTRY\_T2 Normalised CPU time (kSI2K) by COUNTRY\_T2 and VO.  
LHC VOs (and Other VOs). January 2011 - November 2011.

The following table shows the distribution of Normalised CPU time (kSI2K) grouped by COUNTRY\_T2 and VO (only information about LHC VOs is showed in detail. The rest of VOs will be grouped in a new category).

Normalised CPU time [units: 1K, SI2K Hours] by COUNTRY_T2 and VO							
COUNTRY_T2	alice	atlas	cms	lhcb	Other VOs	Total	%
Australia	0	2,408,045	0	0	62	2,408,107	0.19%
Austria	0	1,890,313	2,380,909	0	10,663	4,281,885	0.34%
Belgium	0	0	5,636,182	0	929,737	6,565,919	0.52%
Brazil	73	45	663,337	0	0	663,455	0.05%

- The layout of graphs and the colour schemes have changed
- This was a complete internal rewrite so watch out for changes in behaviour and report them,
- More flexible graphics
  - Bugs have been raised about the labelling on the graphs.
- Query interface returns XML
  - Allows saving query and programmatic interface
  - See the Custom\_view – you can save the link
  - Needs documenting

# EGI ACCOUNTING PORTAL



GLOBAL View

VO MANAGER View

VO MEMBER View

SITE ADMIN View

USER View

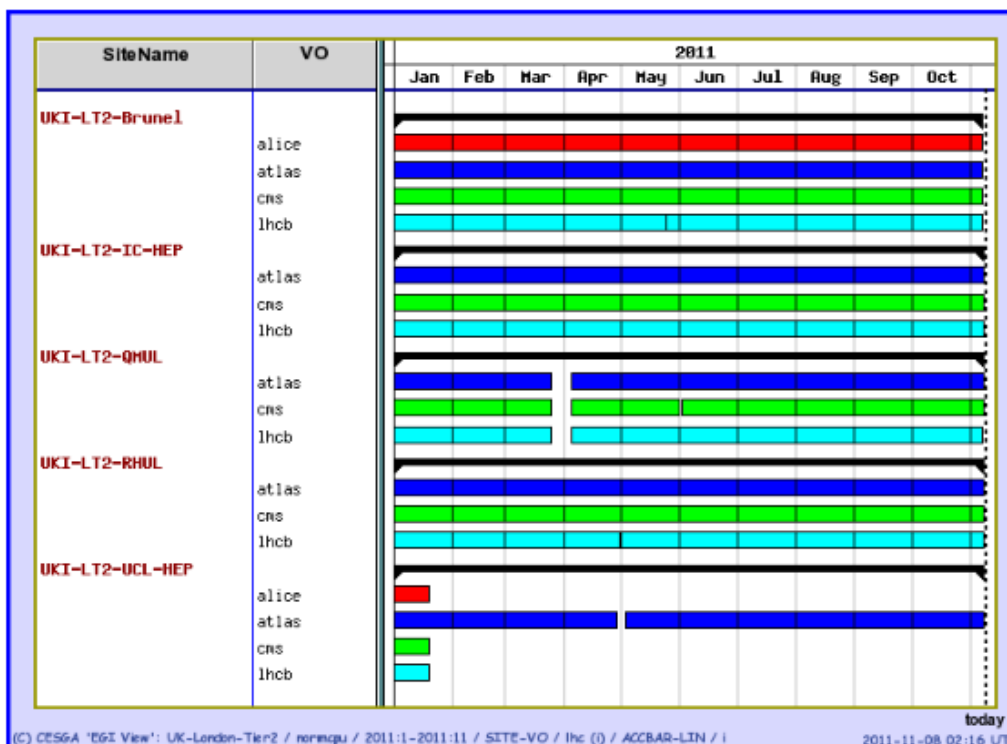
REPORTS

METRICS PORTAL

LINKS

- ▶ Portugal
- ▶ Republic of Korea
- ▶ Romania
- ▶ Russian Federation
- ▶ Slovenia
- ▶ Spain
- ▶ Sweden
- ▶ Switzerland
- ▶ Taipei
- ▶ Turkey
- ▼ UK
  - ▶ UK-London-Tier2
    - ▶ UKI-LT2-Brunel
    - ▶ UKI-LT2-IC-HEP
    - ▶ UKI-LT2-QMUL
    - ▶ UKI-LT2-RHUL
    - ▶ UKI-LT2-UCL-HEP
  - ▶ UK-NorthGrid
  - ▶ UK-ScotGrid
  - ▶ UK-SouthGrid
- ▶ Ukraine
- ▶ USA
- ▶ Countries
- ▼ EGI
  - ▶ Production
  - ▶ PPS
  - ▶ OSG
  - ▶ EELA
  - ▶ UNREGISTERED
  - ▶ VO\_Discipline
  - ▶ VO\_Metrics
  - ▶ CUSTOM\_view

This chart shows a summary of the accounting records for all sites in the tier2. The records are organised according to VO (only information about LHC VOs is showed in detail. The rest of VOs will be grouped in a new category). A detailed view can be obtained by selecting an individual site.





# EGI ACCOUNTING PORTAL



GLOBAL View

VO MANAGER View

VO MEMBER View

SITE ADMIN View

USER View

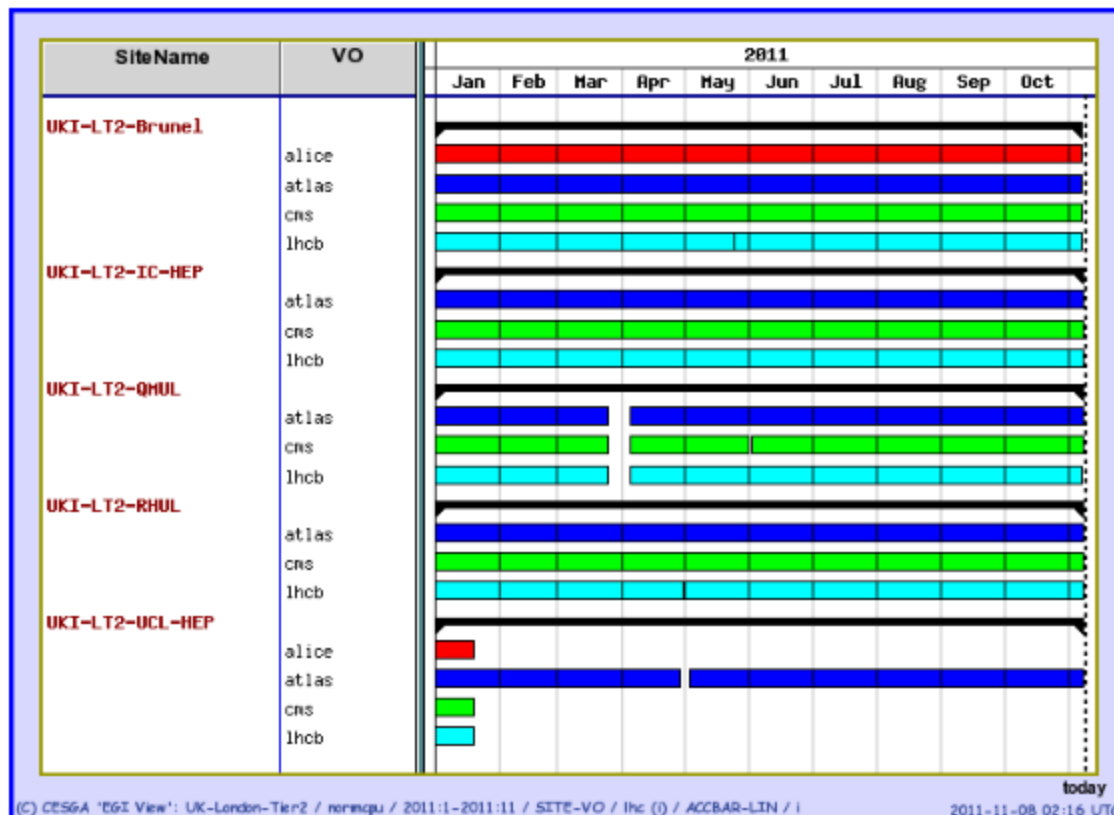
REPORTS

METRICS PORTAL

LINKS

- ▶ Portugal
- ▶ Republic of Korea
- ▶ Romania
- ▶ Russian Federation
- ▶ Slovenia
- ▶ Spain
- ▶ Sweden
- ▶ Switzerland
- ▶ Taipei
- ▶ Turkey
- ▼ UK
  - ▶ UK-London-Tier2
    - ▶ UKI-LT2-Brunei
    - ▶ UKI-LT2-IC-HEP
    - ▶ UKI-LT2-QMUL
    - ▶ UKI-LT2-RHUL
    - ▶ UKI-LT2-UCL-HEP
  - ▶ UK-NorthGrid
  - ▶ UK-ScotGrid
  - ▶ UK-SouthGrid
- ▶ Ukraine
- ▶ USA
- ▶ Countries
- ▼ EGI
  - ▶ Production
  - ▶ PPS
  - ▶ OSG
  - ▶ EELA
  - ▶ UNREGISTERED
  - ▶ VO\_Discipline
  - ▶ VO\_Metrics
  - ▶ CUSTOM\_view

This chart shows a summary of the accounting records for all sites in the TIER2. The records are organised according to VO (only information about LHC VOs is showed in detail. The rest of VOs will be grouped in a new category). A detailed view can be obtained by selecting an individual site.




Normalised CPU time (kSI2K) by TIER1 and VO

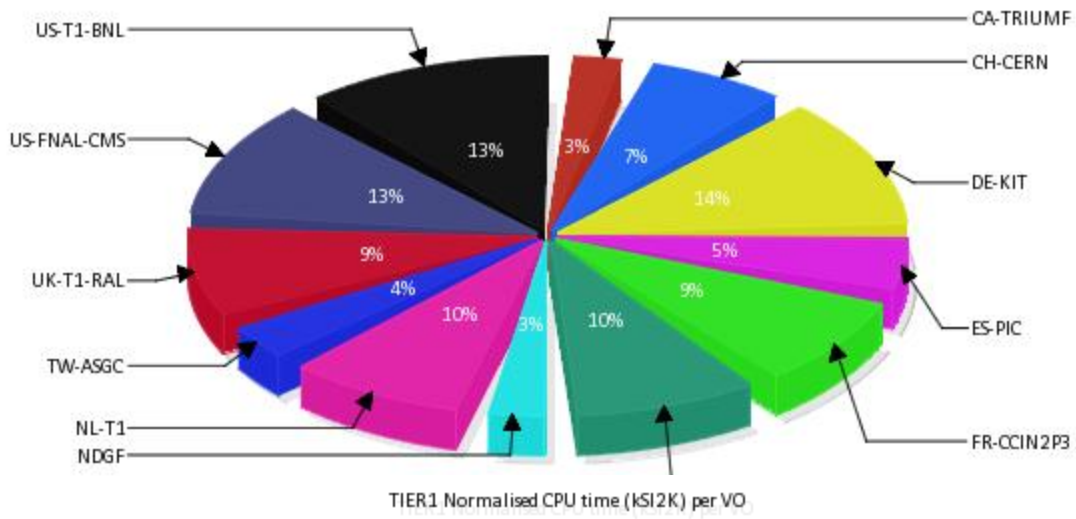
- CA-TRIUMF
- CH-CERN
- DE-KIT
- ES-PIC
- FR-CCIN2P3




Normalised CPU time (kSI2K) by VO and DATE

- alice
- atlas
- cms
- lhcb





© CESGA EGI View: TIER1 / normcpu / 2011:1-2011:11 / VO-



- Accounting records for cpu used by batch jobs are published in several ways to the central APEL Repository.
- This is the worldwide repository for the LHC VOs.
- It publishes summaries of the data to the Accounting Portal @ CESSGA where it can be dynamically queried and a number of standard reports produced.

# Can View:

- Njobs, cpu, wallclock, normalisedcpu, normalisedwallclock, (all in HS06 or kSI2K) cpuefficiency,
- As a table, graphs, and download
- Tables and graphs of one of above, download of all.
- Any timespan of months.
- Any VO, all, or a selection
- Trees for Countries, NGIs, T1, T2
- Display any point in the tree and below from top to site level.
- Date, VO, (Site/Region/Country) as a function of one of the others.

- Reports:
  - T1
  - T2 including pledges
- Views
  - User – see your own jobs
  - VO Member – see FQAN in VO (top 10)
  - VO Manager – see UserDN in VO (top 10)
  - Site Admin – see UserDN at site, FQAN just added.

- The portal will be developed to view the new types of accounting
- What else would you like to see in the portal?

- Tiziana Ferrari and Peter Solagno of EGI are leading a task force to consider all aspects of accounting, bringing together all stakeholders.
- WLCG should be represented



- New APEL Infrastructure
  - next six months
- New Accounting Portal
- New Types of Accounting
- New Requirements?