




Enabling Grids for
E-science in Europe

Accounting in LCG



Dave Kant

CCLRC, e-Science Centre

APEL in LCG/EGEE

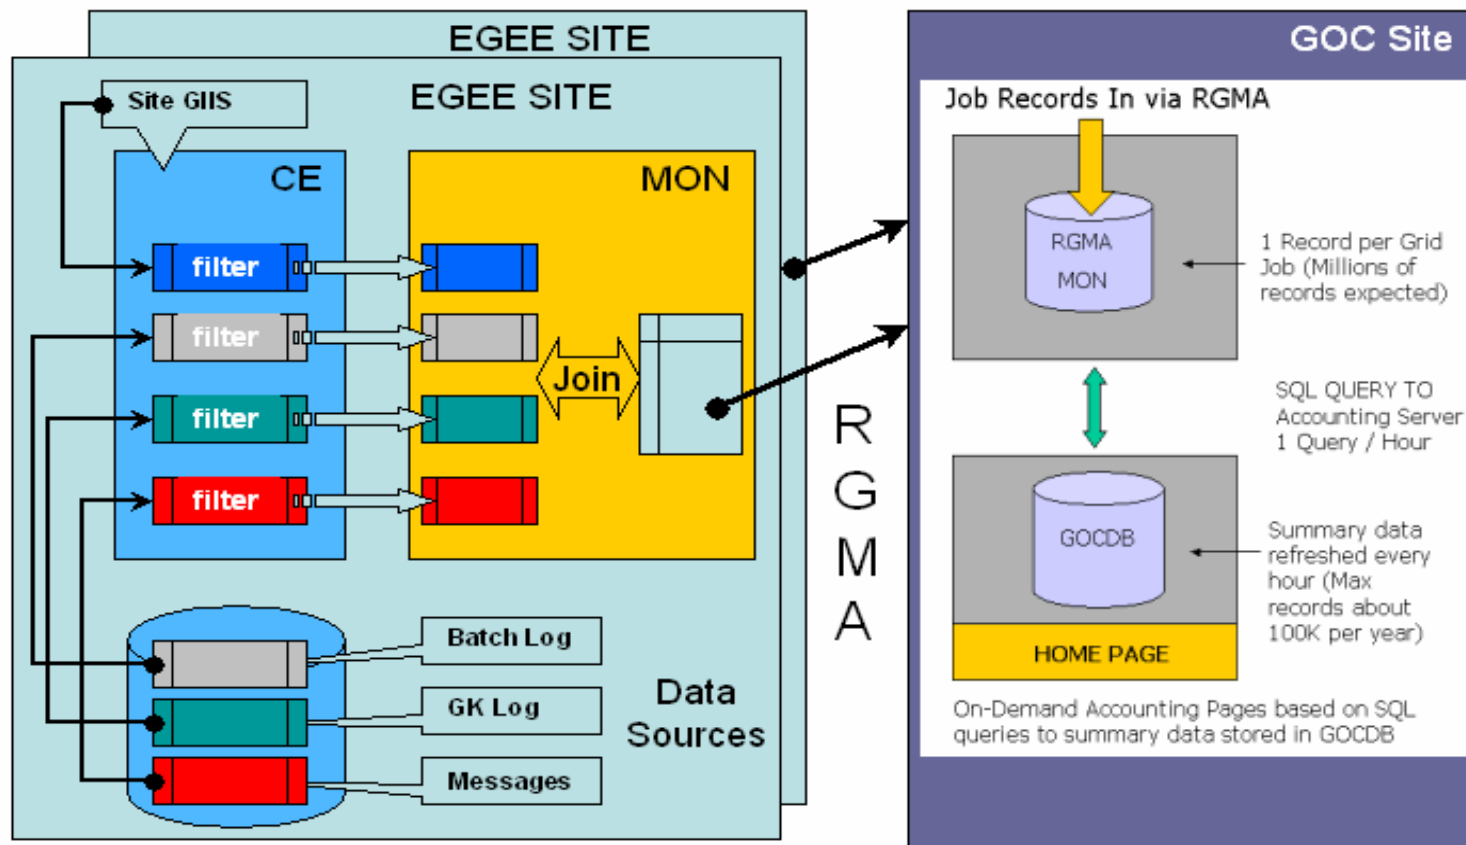


1. Overview
2. Batch System Support
3. Reporting

Accounting Flow Diagram

Batch Supported (PBS/Torque, LSF)


Under Development (SGE, Condor)




Reporting Web Pages

GOC

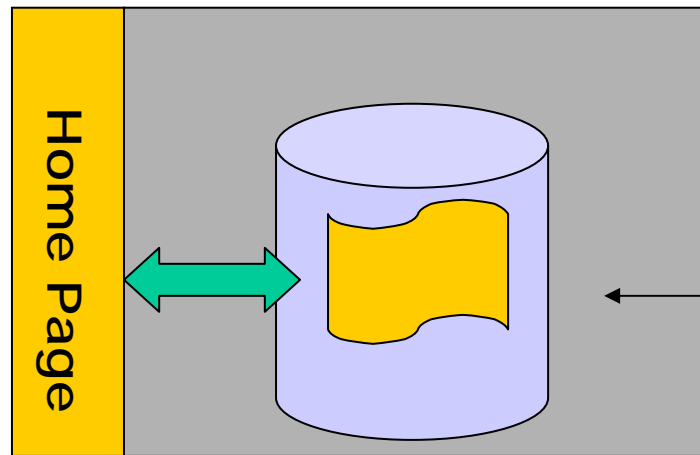
Job Records In via RGMA



1 Record per Grid Job (Millions of records expected)



SQL QUERY TO
Accounting Server
1 Query / Hour



Summary data
refreshed every
hour (Max
records about
100K per year)

On-Demand Accounting Pages based on SQL queries to summary data



User queries

Graphs

Accounting

> Accounting Home

Accounting Plots

> CIC View

> ROC View

> Country View

> Custom Query

> LHC View

> Privacy Statement

General

> LSF Deployment

> Sending Data to us

GOC Accounting Services

Latest News

- August-1-2005:
New version released for LCG 2.6




Accounting RPMs Release-Notes Installation Guide Schema FAQ

<http://goc.grid-support.ac.uk/>


Database Statistics

Number of CEs



■ Distinct Sites = 107 entries
Last Build: Sep 02 12:03:26 2005

Number of Job Records



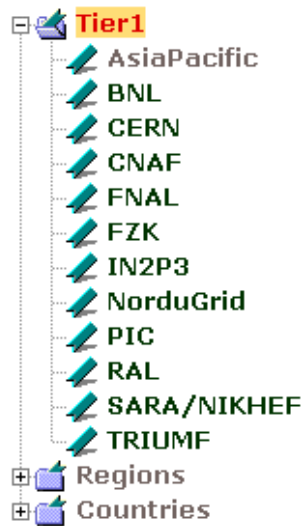
■ Record Count = 3304733 entries
Last Build: Sep 02 12:03:26 2005

107 Sites publishing data (Sep 02 2005)



Over 3.3 Million Job records

~ 100K records per week (period June 1st – mid Aug 2005)

LHC Hierarchical Tree



Reporting: Summary Records provide anonymous global views of data.



Each Tier1, per VO, per Month

Group / Site	VO	Jan 2004	Feb 2004	Mar 2004	Apr 2004
Italy	2688	0	0	14619	0
	alice	0	5040	220333463	0
	atlas	0	3	877	0
	babar	0	0	0	0
	cms	0	0	5776	0
	dteam	0	1560	9208	0
	dzero	0	0	0	0
	lhcb	0	0	47	0

Group / Site	VO	Jan 2004	Feb 2004	Mar 2004	Apr 2004
UK / Ireland	2688	0	0	0	0
	alice	0	410303	94319599	589065
	atlas	4	0	1408	76006
	babar	25	0	7	3
	cms	0	47	1341	1137
	dteam	522	11719	24667	26962
	dzero	0	0	28	0
	lhcb	0	0	64	141

Other
Distributions
Normalised
CPU
Jobs

Each Site, per VO, per Month

Group / Site	VO	Jan 2004	Feb 2004	Mar 2004	Apr 2004
CNAF	2688	0	0	14619	0
	alice	0	5040	220333463	0
	atlas	0	3	877	0
	babar	0	0	0	0
	cms	0	0	5776	0
	dteam	0	1560	9208	0
	dzero	0	0	0	0
	lhcb	0	0	47	0
	INFNLNF	2688	0	0	0
alice		0	0	0	0
atlas		0	0	0	0
babar		0	0	0	0
cms		0	0	0	0
dteam		0	0	0	0
dzero		0	0	0	0
lhcb		0	0	0	0
LEGNARO		2688	0	0	0
	alice	0	0	0	0
	atlas	0	0	0	0
	babar	0	0	0	0
	cms	0	0	0	0
	dteam	0	0	0	0
	dzero	0	0	0	0
	lhcb	0	0	0	0

Accounting Overview

1. Overview
2. Challenges Ahead
3. World Wide Accounting Service for LCG

Challenges Ahead

- Recognise that accounting isn't just about "job usage" its about Resource usage which encompasses many things:-
 - CPU Usage ☺
 - Also Storage & Network Usage (Who should we talk to?) ☹
- How do we describe this data?
 - Luckily there is a GGF Usage record which provides a generic description of resource usage ☺
 - Are these descriptors stable?
 - Are they sufficient to describe the data?
- How important is accounting?
 - Compute resource viewed as a grid currency
 - Need a guarantee that the data has not been tampered with in an un fair way
 - How does normalisation fit into this? The concept of a raw usage records has no meaning if internal scaling is applied to Heterogeneous farms.
 - GGF UR allows a "cost" descriptor
 - Do we need an agreement of cost?

Challenges Ahead

- Data Collection
 - Many implementations for collecting accounting data in LCG World
 - APEL/DGAS in EGEE
 - SGAS in SweGrid
 - OSG (See talk by Mattio)
 - Sites that implement their own systems (Fermilab: multiple grid job managers from different grids feed a single condor pool)
- No mechanism in place to share this data in a consistent way in place.
 - GGF Working on a Resource Usage Service ☺
 - What would the model for data sharing look like? Low level or high level?
 - Low Level: sensors publishing data via a web service?
 - **High level: Data collected within the infrastructure, aggregated in a meaningful way, reviewed and approve data before it can be passed on (FermiLab)**
 - Some Tier-1 centres have concerns about data association
“LCG not EGEE” “Will the service be separate?”

Challenges Ahead (This Slide Highly Recommended by JT)

- Usage Reporting at what Level?
 - **Anonymous level: How much resource has been provided to each VO**
 - Aggregation across: VOs, Countries, Regions, Grids, Organisations
 - Granularity: summed over units of Hours, Days, Weeks, Months?
- User Level Reporting?
 - **If 10,000 CPU hours were consumed by Atlas VO, who are the users that submitted the work?**
 - Data privacy laws
 - A Grid “DN” is personal information which could be used to target an individual.
 - Who has access to this data and how do you get it?
 - Can CA policies change to support anonymous DNs and reverse DN mappings?
 - What are the consequences? Are there any lawyers in the audience?

World Wide Accounting Service for LCG

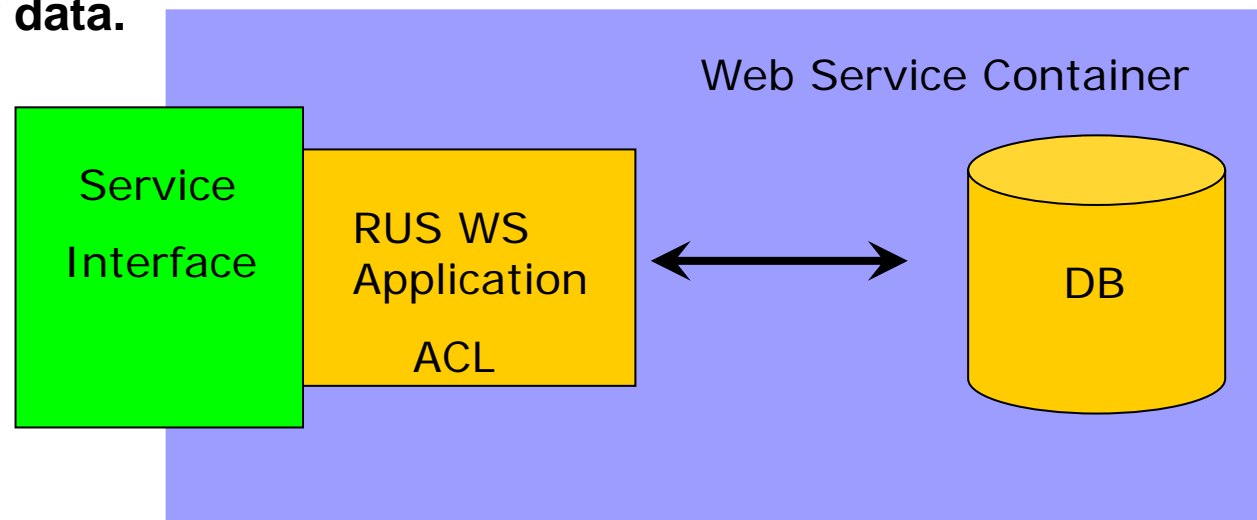


- ***Project involves combining results from all three peer infrastructures and presenting an aggregated view of resource usage for LHC VOs to the RRB***
 - ***Peer Infrastructures in LCG***
 - *Open Science Grid + Others (Ruth Pordes, Philippe Canal, Matteo Melani)*
 - *Nordugrid (Per Oster, Thomas Sandholm)*
 - *LCG/EGEE (Kors Bos, Dave Kant)*

GRID-ACCOUNTING@LISTSERV.RL.AC.UK

Resource Usage Service

- Based on emerging GGF standards and Web Services
 - GGF UR, OGSF
- An implementation exists in “Market for Computational Science” – UK e-Science project
- Use case might be:
 - A user invokes the query service through a web browser, using SSL for client authentication, to ensure that usage information at user level belongs to the user. Servlet sends query to RUS web service and gets user data.



Possible Roadmap

- **Stage 1:** Lets try to get some data from each of Tier-1s
summary records describing VO usage over a finite period of time
 - Before end 2005
- **Stage 2:** Centralised database with a web service interface (RUS)
to publish/query accounting data (summary records)
 - Sometime in 2006
- **Stage 3:** Distributed databases with a complete RUS
implementation including permission model.
 - Sometime early 2007

--- END ---