




# Accounting @ User-Level, VOMS Groups and Roles

Dave Kant

CCLRC, e-Science Centre



# Outline

- User-Level Accounting
- Group and Role Accounting
- Accounting Portal
- Storage Accounting
- Summary



# User-Level Accounting

- **LCG JSPG User-Level Accounting Data Policy**
  - Discusses the treatment of accounting data
  - Topics covered
    - User Consent, What is stored and where? Who has Access rights? For what purposes? Confidentiality
  - Document in progress
- **APEL Feature**
  - Sites can switch on User DN encryption when publishing data
  - 1024-bit RSA public/private key-pair to maintain confidentiality
- **Request:**
  - Policy document isn't official yet; we ask the Tier-1s publish encrypted User DNs.
  - 18 sites have already done this:-  
CESGA-EGEE , CIEMAT-LCG2, CNB-LCG2, CSCS-LCG2, **FZK-LCG2**,  
IFAE, IFCA-LCG2, INFN-TORINO, ITEP, JINR-LCG2, LIP-Lisbon,  
**RAL-LCG2**, RU-SPbSU, TAU-LCG2, UB-LCG2,  
UKI-NORTHGRID-LANCS-HEP, UPV-GRyCAP , USC-LCG2



# Group and Role Accounting

- Requirement
  - Define how CPU resources should be allocated to different activities
    - 60% of CPU to production at site X for the next three months
  - Measure what has been delivered
    - In last quarter of 2006, site X delivered Y% production to VO\_1
- What needs to be in place?
  - Users should get a VOMS proxy before job submission
  - Extract the users VOMS proxy from the CE
  - Include this information in the job accounting record
- Implementation
  - On LCG-CE via the Grid Accounting Mapping File (Patch #898)
    - Maps grid credentials to local batch resources
      - Users FQAN Chain
      - LCMAPS will map the user according to the VOMS FQANs in the proxy
  - APEL Patch #1047 in Certification
    - Tests performed against Patch #898 (CERN, CNAF)
    - We take the whole chain (can be quite long)
    - Group and Role determined from **PRIMARY** part of UserFQAN chain



# What FQANs have been published?

- Activity Related

- Breakdown of usage according to activity and VO

```
/cms/Role=production/Capability=NULL
```


```
/lhcb/lcgprod/Role=NULL/Capability=NULL
```

```
/atlas/Role=lcgadmin/Capability=NULL
```

- Chain with primary and secondary components.

- No activity information in the primary
- Can only account usage to the VO

```
/dteam/Role=NULL/Capability=NULL;/dteam/cern/Role=NULL/Capability=NULL
```



# Test Record

- What you see in R-GMA

**GlobalJobID:**

`https://lxb1762.cern.ch:9000/-cfWmtwoGI5CUwGrI5gVmA`

**UserDN:**

`*APEL V.0.2*`

`P7IMQIkPndWbCpeecX/iqqBhqYdsOp+DZl+9+o9GQ3BPXx8uzcHUqZOJQsq8F6KIczAvxSw+I8x8  
lwHJ6l9kxIyYbLTEkELI3Ul77I6zhzT90zUDLEpgsQ+0XY3tPRGwu5uG/ibtcWxefOtvZoM6FWwf  
8yZmv8yLpjmNkMxZotI=`

**UserFQAN:**

`/dteam/Role=NULL/Capability=NULL;/dteam/cern/Role=NULL/Capability=NULL;`

**ExecutingCE:**

`lxb2034.cern.ch:2119/jobmanager-lcgpbs-dteam`



# Is User FQAN Confidential?

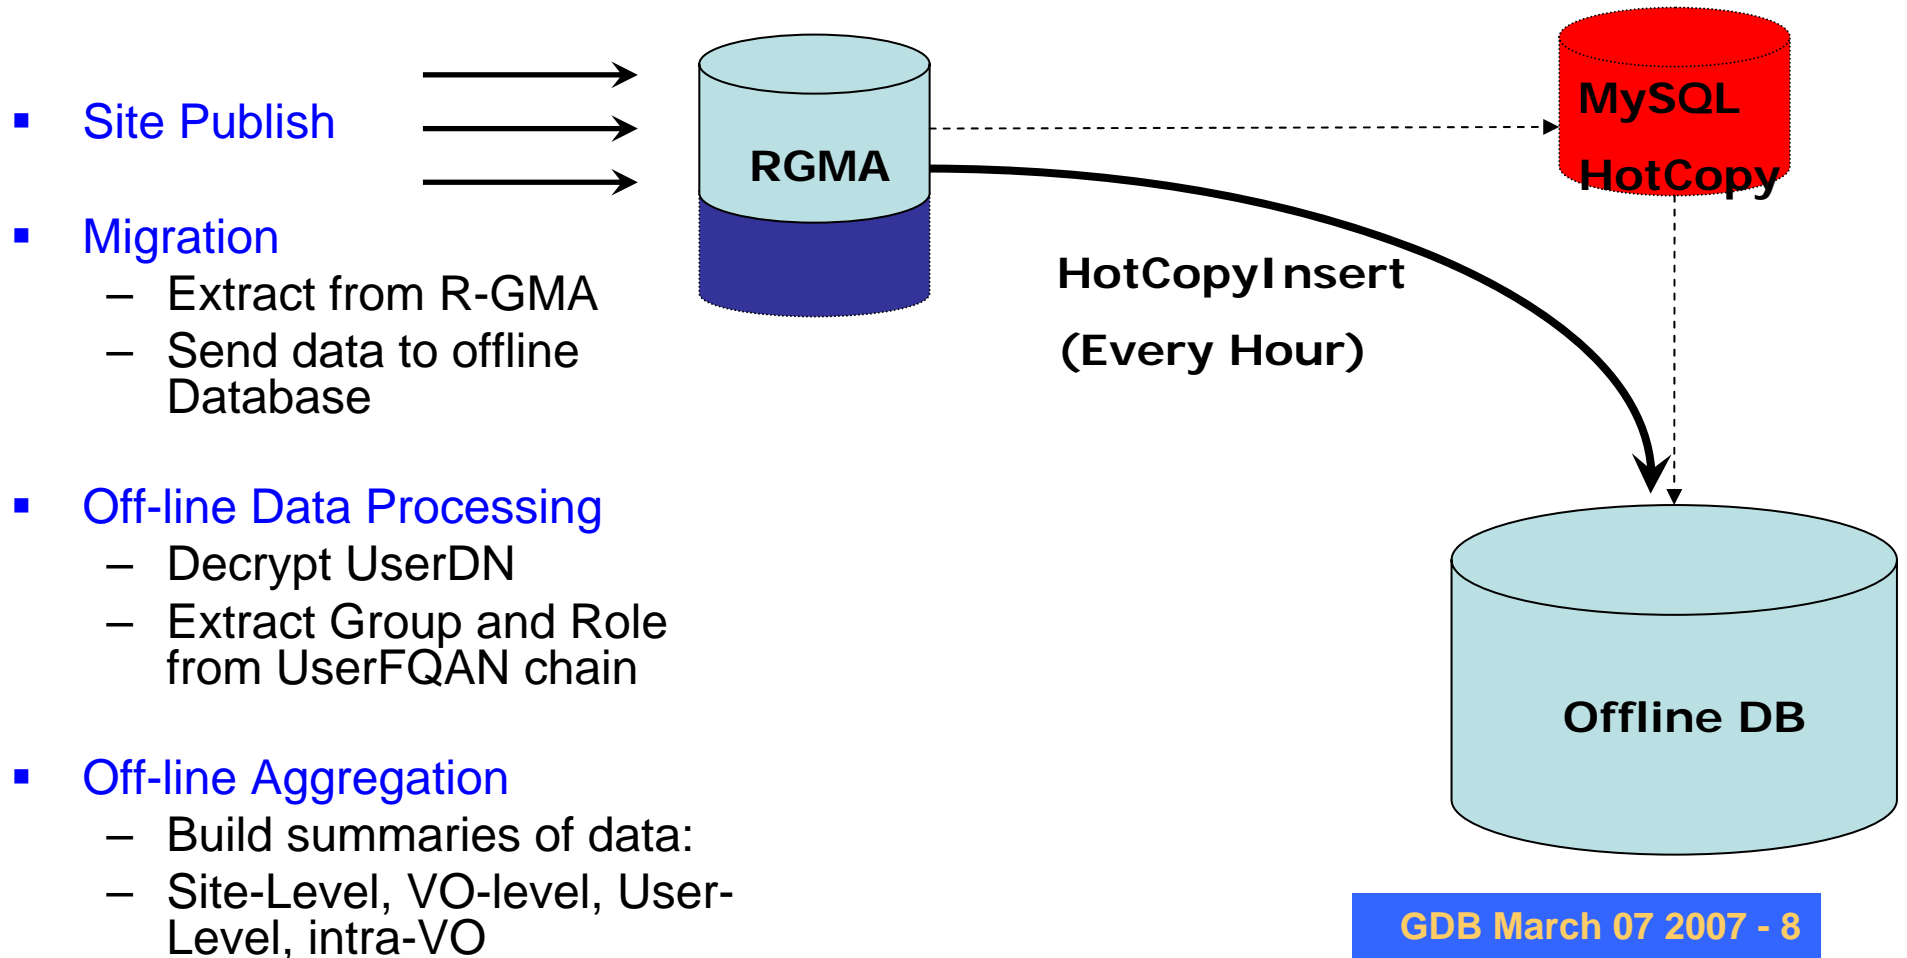
- User FQAN Chain readable in R-GMA by anyone with an IGTf approved certificate.
- Should we encrypt this information?
- If yes, can explore two options.
  - Encrypt only the primary FQAN
    - Throw away the secondary/tertiary parts ...
    - Easy to implement.
  - Encrypt the whole chain
    - More complicated , because the FQAN Chain length can be very long and encryption is limited by the length of the public key.
    - RGMA schema change results in accounting data in two archivers ...



# Data Life Cycle

## From Site to Portal

Desirable to refresh the data shown on the portal at frequent intervals  
Aim for hourly updates







# Portal

## ■ Accounting Portal

- Access Right according to the “Actor Model”
- Five Actors:
  - Users, VO-Resources Manager, VO-Member, Site-Admin
  - GOC Admin
  - Three are implemented via DN proxy and ACLs
  - VO-Member access requires the VOMS proxy information.
    - How do we do this via the web?

## ■ Display Features

- Build Resource statistics as a function of User and FQAN
- How much CPU consumed by User X at Site Y
- Top Ten Users in a VO
- How much production work done by VO X at Site Y



# VO Resource Manager

- Table shows CPU, WCT and Job Eff. of the Top 10 Anonymised Users
- Breakdown of Usage: DN / VO / Group / Role

EGEE View

**VO MANAGER View**

VO MEMBER View

SITE ADMIN View

USER View

January 2006 - December 2006.


The following table shows the Usage of the Top 10 Users ordered by Normalised CPU time and the Total Usage of the Other Users. A detailed view can be obtained by selecting an individual user.

Top 10 Users ordered by Normalised CPU time													
User		Jobs		CPU time		Norm. CPU time		WCT		Norm. WCT		CPU Efficiency	Avg. CPU time
#	ID	#	%	Hrs	%	Hrs	%	Hrs	%	Hrs	%	%	Hrs
1	<a href="#">2daf892f6a1f68d0</a>	15,629	25.2%	120	24.3%	88	22.6%	1,314	16.5%	990	16.0%	9.1	0.01
2	<a href="#">671e5b277fa6d1ac</a>	20	0.0%	98	19.8%	87	22.3%	103	1.3%	91	1.5%	95.1	4.90
3	<a href="#">4cd78d6b029f7050</a>	7,773	12.5%	80	16.2%	55	14.1%	924	11.6%	765	12.4%	8.7	0.01
4	<a href="#">5d61bd2201bec5f9</a>	7,816	12.6%	52	10.5%	41	10.5%	331	4.2%	250	4.0%	15.7	0.01
5	<a href="#">2a153e141e98f06a</a>	1,950	3.1%	36	7.3%	36	9.2%	45	0.6%	46	0.7%	80.0	0.02
6	<a href="#">62d3866c3c8260d6</a>	39	0.1%	27	5.5%	16	4.1%	31	0.4%	18	0.3%	87.1	0.69
7	<a href="#">2756dfcb65975a47</a>	95	0.2%	19	3.8%	15	3.8%	22	0.3%	18	0.3%	86.4	0.20
8	<a href="#">748206ea352cce31</a>	467	0.8%	12	2.4%	14	3.6%	15	0.2%	16	0.3%	80.0	0.03
9	<a href="#">2d04c1be5d64c1b8</a>	3	0.0%	11	2.2%	13	3.3%	26	0.3%	21	0.3%	42.3	3.67
10	<a href="#">2d72edb26620a697</a>	83	0.1%	9	1.8%	6	1.5%	125	1.6%	79	1.3%	7.2	0.11
Others (DN known)		25,073	40.5%	18	3.6%	10	2.6%	4,695	58.9%	3,527	57.1%	0.4	0.00
Others (DN unknown)		3,021	4.9%	12	2.4%	9	2.3%	343	4.3%	356	5.8%	3.5	0.00
<b>Total</b>		<b>61,969</b>		<b>494</b>		<b>390</b>		<b>7,974</b>		<b>6,177</b>		<b>6.2</b>	<b>0.01</b>

[Click here for a csv dump of this table](#)

Key: 0% <= eff < 75%; 75% <= eff < 90%; 90% <= eff < 100%; eff >= 100% (parallel jobs)

GDB March 07 2007 - 10



# Other Accounting Issues/Status

- **OSG**

- Deployment of Gratia accounting system across the CMS and ATLAS tier 1 and 2 centers.
- Expect monthly summaries.

- **DGAS**

- Validating sites after the deployment of Patch #898
- Publish job records from T1 and T2's daily

- **Comparison between CERN and APEL**

- Found a large discrepancies in usage numbers which was traced to a bug in multiple CE support in APEL.
- Implemented a bug fix.
- Repeat the test for the February 2007 dataset.



# Storage Accounting

- Display of Storage Information published by GIPs into the information system.
- Query top-level BDII, write data to MySQL, process the data and build RRD graphs
- Used and Allocated, Disk and Tape at the VO level
- Visualisation of Storage Used and Allocated per VO for Disk and Tape
- No breakdown of these numbers at a deeper level
  - Group/Role
  - File Transactions (Uploads, downloads)
  - Protocols



<http://goc02.grid-support.ac.uk/accountingDisplay/view.php>

Select Resources via a Tree

Select time interval (last year, last month, last week, last day)

Select VOs, SEArchitecture

### EGEE Hierarchical Tree

- Production
  - AsiaPacific
  - CentralEurope
  - CERN
  - France
  - GermanySwitzerland
  - Italy
  - NorthernEurope
  - Russia
  - SouthEasternEurope
  - SouthWesternEurope
- UKI
  - GridIreland
  - LondonT2
  - NGS
  - NorthGrid
  - ScotGrid
  - SouthGrid
  - Tier1A
- PPS

### Storage Accounting Display (Version 0.3)

**Select Interval:**  
last day

**VO Groups**     LHC     non-LHC     ALL     Custom

alice    atlas    babar    biomed    calice  
 cdf    cedar    cms    dteam    dzero

**VOs:**  
 esr    fusion    geant4    hone    ilc  
 lhcb    magic    minos    na48    ngs  
 ops    oxg    pheno    planck    ralpp  
 sixt    t2k    zeus

**SEArchitecture:**    disk    tape    unknownArch



## CSV dump of current Used and Allocated storage

### Disk Usage Only (TB)

Disk Used and Allocated per VO (Units are TB)							
Org Unit	Site	VO				Total (TB)	Percentage
SouthGrid		alice	atlas	cms	lhcb		
	EFDA-JET	0.190 <b>1.570</b>	0.190 <b>1.570</b>	0.190 <b>1.570</b>	0.190 <b>1.570</b>	0.76 <b>6.28</b>	0% <b>12.5%</b>
	UKI-SOUTHGRID-BHAM-HEP	0.000 <b>0.000</b>	1.379 <b>1.379</b>	0.001 <b>0.001</b>	0.000 <b>0.000</b>	1.38 <b>1.38</b>	0% <b>2.7%</b>
	UKI-SOUTHGRID-BHAM-PPS	0 <b>0</b>	0 <b>0</b>	0 <b>0</b>	0 <b>0</b>	0 <b>0</b>	0% <b>0%</b>
	UKI-SOUTHGRID-BRIS-HEP	0.000 <b>0.785</b>	0.028 <b>0.812</b>	0.824 <b>2.360</b>	0.000 <b>0.785</b>	0.852 <b>4.742</b>	0% <b>9.4%</b>
	UKI-SOUTHGRID-CAM-HEP	2.306 <b>3.220</b>	2.306 <b>3.220</b>	2.306 <b>3.220</b>	2.306 <b>3.220</b>	9.224 <b>12.88</b>	0% <b>25.6%</b>
	UKI-SOUTHGRID-OX-HEP	0.000 <b>0.861</b>	1.955 <b>2.816</b>	0.001 <b>0.862</b>	0.000 <b>0.861</b>	1.956 <b>5.4</b>	0% <b>10.7%</b>
	UKI-SOUTHGRID-RALPP	0.000 <b>0.131</b>	0.716 <b>4.194</b>	11.160 <b>14.680</b>	0.000 <b>0.655</b>	11.876 <b>19.66</b>	0% <b>39.1%</b>
Table Summaries							
	Used (TB)	2.496	6.574	14.482	2.496	26.048	
	<b>Allocated (TB)</b>	<b>6.567</b>	<b>13.991</b>	<b>22.693</b>	<b>7.091</b>	<b>50.342</b>	
	Percentage	9.6	25.2	55.6	9.6	100	
		<b>13%</b>	<b>27.8%</b>	<b>45.1%</b>	<b>14.1%</b>	<b>100%</b>	




CSV dump of this table



# Other Storage Accounting Plans and Issues

- Enhance the LHC view of storage resources
  - Tier-1 tree
- Extend breakdown per site to the RRD graphs
- To show AverageUsage, AverageAllocated per Month per VO per LHC Tier1.
  - Example: RaI-LCG2-ALICE use 50TB tape in the first week of Jan, and then deleted the data.
- How do we get data from OSG?
- Bring the Storage and CPU reporting together into a single Portal.



# Summary

- User level accounting is deployed and many sites are publishing. T1s are encouraged to deploy.
- This is not sufficient for the role/group based accounting that VOs have requested.
- New release of APEL in certification; together with the accounting mapping log file, FQAN accounting is possible.
- Storage Accounting Visualisation tools displaying Used and Allocated per VO; LHC view is needed