Accounting Policy Update

John Gordon/Dave Kelsey
CCLRC-RAL

LCG Grid Deployment Board
June 2006

History

- Dave presented a draft outline on a User-Level Accounting Data Policy at GDB in Rome, April 2006
- There has been a little progress since then

Reminder

- Dave proposed that the following issues needed to be covered in such a policy:-
- User Consent
  - In the Grid AUP the user consents to accounts being stored.
- What is stored and where?
  - log files and accounting records derived from them
  - user identification (DN, FQAN)
- Access rights - Who can read?
  - authorised people including the user
- Access for what purposes?
  - accounting only, keep confidential

Holger Martin of FZK legal advise.

- 1. Schema proposed does not involve any *sensitive* user information (dates of birth, addresses, phone numbers etc.), thus allowing for more light weight data handling and protection.
- a. *Anonymous, statistical* accounting information of type "aggregated consumed CPU time per VO / site / month".
  - This type of data is mostly uncritical and may be provided in a "world readable" way to scientific boards and communities.
- b. *User-related* accounting information.
  - This type of data can potentially be used to record and control the work of individual persons, and allows conclusions about his/her working methods, results, performance etc.
  - This must be prohibited.

- 3. Accounting requirements. The policy must contain a list of strong arguments for the necessity of collecting accounting data and the purpose thereof
- 4. Access to the (user-level) data. The current policy envisions to provide data access for two different entities: authorized GOC and VO managers.
- These two groups of persons must belong to states that accept the European laws for data privacy (i.e., that user-level data are only exchanged within such states).
  - Currently, these are the 25 EU member states, the EEA member states and several other states (including Canada and Switzerland) protection.

- Exchange of private data with organisations in the U.S. is documented in a special treaty called "safe harbour privacy principles", and the organisation to receive / work with private data should verify to accept these principles.
- Those people *handling* user-level data (currently the GOC and VO managers), that they sign the policy i.e. especially use the data exclusively for the described purposes and don't provide them to non-authorized persons (i.e., don't misuse them).

- 6. Routine usage of the data. **The policy should describe the routine usage of the data as completely as possible.**
- 7. Misuse of and suspicion on wrong accounting data. **The policy should describe a procedure to be followed in such cases.**
- **It would be useful to get feedback from other countries although this advise was given in the international context.**

Proposals

- Sufficient to encrypt the DN in transit
  - a single record is low risk and the information is not private
- Anonymise DN in database
  - can see behaviour of users without knowing their identity
- Strictly control access to DN identities

Who needs to see user level records?

- The VO is the group of people who are members of the VO.
- Members are added or deleted to/from the VO by the VO administrator. The VO Manager (aka VOMS admin) is the only person that has the appropriate rights to modify this list. The decision to add or delete a member to list is with the VO Manager who will have to follow the recommendations of the management team of the collaboration.

- The VO Resource Manager is a new role we are proposing. They are the only person who has access to all data in the accounting database belonging to her VO including the DN. It is the responsibility of the VO Resource Manager to use the DN related information appropriately and make sure this information does not proliferate beyond the circles where it is needed. People with the 'role' of VO Manager must be given access to all accounting records in the database for their VO except for the DN information.
- GOC Developers should also sign the policy
- The User has a right to see own records

Do the VOs agree?

- Does this concept of VO Manager and VO Resource Manager fit well with existing VO practice?
  - better names?
- The VO Resource Manager could be a role given to multiple people but it should be restricted. Not just 'anyone who needs to see user data'

Further Thoughts

Requirements for an LCG Accounting System:

- Must preserve all accounting records for the LHC VOs
- In the transfer of accounting records the DN must be encrypted
- Access to the accounting data must be protected
- Access to CPU/VO allowed to anyone from that VO
- Also access to Group/Role/DN is only for VO management
  - *This needs to be better defined, as above …*
- Aggregated CPU usage per VO may be published on the web
- Aggregated CPU usage per Group/Role only available to VO
- Aggregated CPU usage per DN only available to VO managers
- Individual job records must be deleted after 1 year
- Aggregated data may be kept for the duration of the project

Next Steps

- Work on the policy will continue within ?????
- Anyone holding user level information should be interested in this.
  - including all the various monitoring and logging systems
- See also Accounting Policies Paper on Agenda.