JSPG
Status and plans

EGEE'06 Conference
Geneva, 28 Sep 2006

David Kelsey
CCLRC/RAL
d.p.kelsey@rl.ac.uk

www.eu-egee.org

Information Society and Media

Overview

- **Introduction to JSPG**
- **Interoperable policies**
- **Current policy document set**
- **Recently approved documents**
- **Current work**
  - New top-level Security Policy
  - Site Operational Procedures Policy
  - User-level Accounting and Data Privacy
- **Future plans**

Introduction to JSPG

- **Joint Security Policy Group**
- **Mandate**
  - To advise and make recommendations to the Grid Deployment Manager and GDB on all matters related to Security
    - Policies are agreed and adopted by GDB/MB for WLCG
    - Policies are agreed and adopted by PEB for EGEE
  - To produce and maintain
    - Policies and procedures on Registration, Authentication, Authorization and Security
  - Where necessary recommend the creation of focussed task-forces made-up of appropriate experts
    - E.g. Task force on LCG User Registration

JSPG(2)

- **"Joint" initially means EGEE *and* LCG**
  - Strong participation by USA Open Science Grid
  - Now "Joint" = EGEE/OSG/WLCG and others
- **In EGEE an activity of SA1 (Deployment & Operations)**
  - Discuss all documents with ROC Managers
  - Participation of site managers/security officers
- **Strong links to other security groups**
  - Middleware Security Group
  - Operational Security Coordination Team
  - Grid Security Vulnerability Group
  - EU Grid PMA

JSPG membership

- **Application representatives/VO managers**
  - Discussions with VO managers as/when required
- **Site Security Officers**
  - Bob Cowles (SLAC and OSG), Denise Heagerty (CERN)
- **Site/Resource Managers/Security Contacts**
  - Dave Kelsey (RAL) – Chair
  - Romain Wartel (CERN)
- **Security middleware experts/developers**
  - Joni Hahkala, David Groep, Andrew McNab,
- **CERN Deployment team**
  - Maria Dimou, Ian Neilson (Security Officer)
- **Now expanded to include other EU Grid projects**
  - SEE-Grid, DEISA, Diligent
- **Other EU Infrastructure projects (may) use our policies**
  - BalticGrid, EELA, EUMedGrid, EUChinaGrid, …
- **Volunteers ALWAYS welcome (please contact me!)**

Interoperable Policies

- **Aim to allow applications (VO's) to easily use resources in multiple Grids**
- **The simplest approach**
  - Common Policies
    - User AUP
    - Site AUP
    - VO AUP
    - Operational procedures and other policies
- **If not common then at least not conflicting**
  - Does NOT override local site and network security policy
- **EGEE working with other EU Grid projects**
  - Common policies and procedures
- **EU eInfrastructure Reflection Group (eIRG)**
  - Common approach at highest level
    - EGEE inputs policy for consideration

Security Policy

picture from Ian Neilson

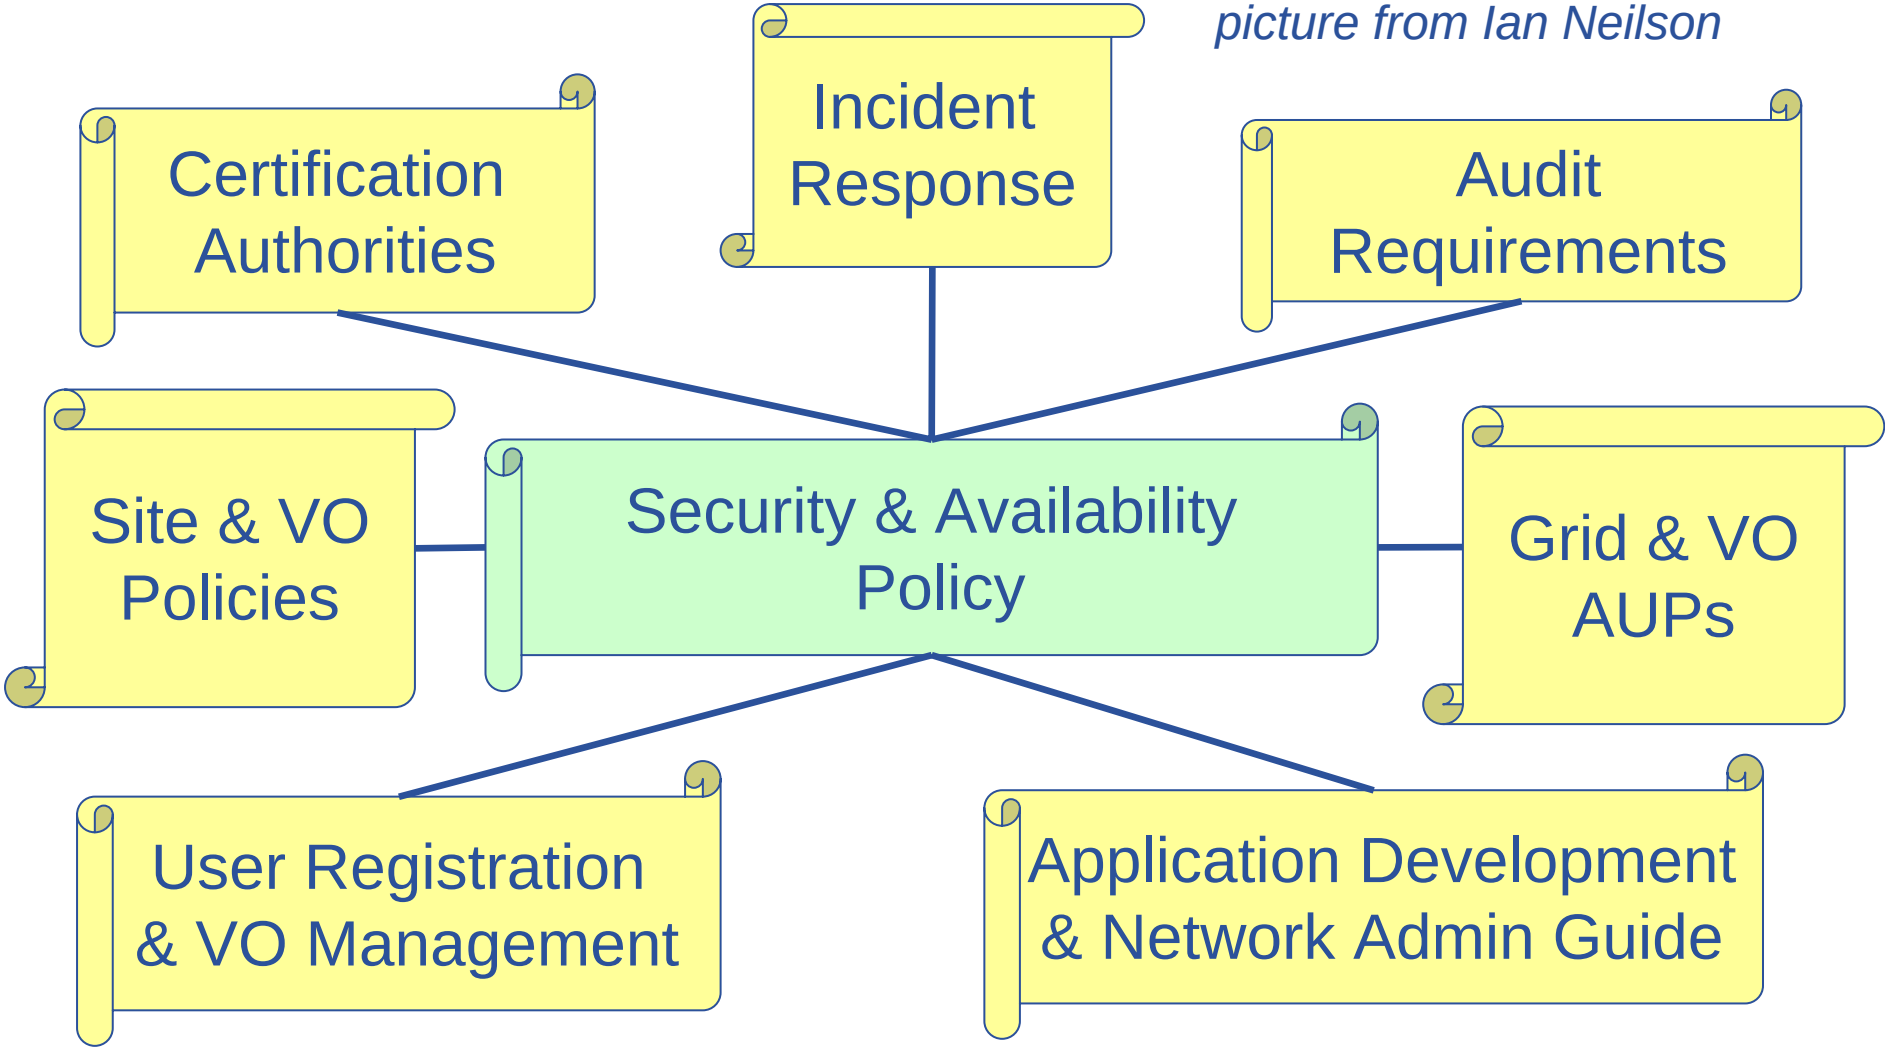

Current policy

The core set

- *Security and Availability Policy for LCG (needs update)*
  - V4c, 17 Oct 2003
- *Grid Acceptable Use Policy*
  - V3.1, 28 Oct 2005
- *LCG/EGEE Virtual Organisation Security Policy*
  - V1.7, 31 Oct 2005

Sub-policy documents

- *Approval of Certification Authorities*
  - V2.5, 3 July 2006
- *Audit Requirements for LCG-1 (needs update)*
  - V1.2, 19 June 2003

Current Policy (2)

Sub-policy documents (continued)

- *Requirements for LCG User Registration and VO Membership Management*
  - V2.7, 1 June 2004
- *Guide to LCG Application, Middleware & Network Security*
  - V1.6, 19 July 2004
- *Site Registration Policy & Procedure*
  - V2.0, 16 Mar 2005
- *LCG/EGEE Incident Handling and Response Guide*
  - V2.1, 15 June 2005

It was agreed by the GDB on 13th Jan 2004

- **all LCG Policy documents remain valid until they are updated**
- **Also adopted by EGEE (and other Grids)**

Obsolete Documents

No longer in use

- *Rules for use of LCG-1 Computing Resources*
  - V2, 23 June 2003
- *User Registration and VO management for LCG-1 in 2003*
  - V1.2, 3 July 2003
- *Approval of LCG-1 Certificate Authorities*
  - V1.0, 2 June 2003
- *LCG Resource Administrators' Guide*
  - V0.2, 16 Feb 2004
- *Procedure for Site Self-Audit*
  - V0.2, 16 Feb 2004
- *LCG Service Level Agreement Guide*
  - V0.3, 16 Feb 2004

Recent JSPG work

- **Recently (last 12 months) approved documents**
  - Grid AUP
  - VO Security Policy (this requires a VO AUP)
  - CA Approval (using IGTF accredited CA's)
- **Several documents need updating**
- **Current work**
  - Top-level Security Policy document (revision)
    - Defines roles and responsibilities, sanctions etc
  - Site Operational Procedures Policy (new document)
  - User-level Accounting data policy (privacy issues)
  - VO Naming (use DNS style) (revision of VO Security policy)

- **All new documents are aimed to be simple and general, e.g. apply to "Grid" not "EGEE" (like the Grid AUP)**

Top-level Security Policy

- **Current draft (V5.1)**
  - http://agenda.cern.ch/askArchive.php?base=agenda&categ=a057709&id=a057709s1t20/document
  - not yet ready for detailed discussion
    - but comments on general approach very welcome!
  - Is it useful to other Grids?
    - If not, what would we have to change?
- **Today, present the structure and approach**

Security Policy

INTRODUCTION

- **Objectives**
  - Regulates activities related to security of the resources
  - Gives authority for actions
  - Places responsibilities on individuals and bodies
- **Scope**
  - Applies to resources, users, administrators, sites, VOs and application developers
  - Does not override local policies
  - Sub-documents (detailed procedures, rules and guides) are part of the policy
- **Ownership, maintenance and review**
  - JSPG

Security Policy (2)

GRID SERVICES and RESOURCES

- **Define the assets that need to be protected**
  - Equipment and software required to run services
  - Data held on these services

ROLES and RESPONSIBILITIES

description of each role and major responsibilities

more details in associated sub-documents

- **Grid Management**
- **VOs**
- **Sites**
- **Resource administrators**
- **Users**
- **Security operations**

Security Policy (3)

- **PHYSICAL SECURITY**
  - Short statement on minimum requirements
- **NETWORK SECURITY**
  - Minimum requirements and firewalls
- **LIMITS TO COMPLIANCE**
  - Local legislation
  - Emergency action and processes for this
- **SANCTIONS**
  - Sites, Users and VOs who fail to comply with policy
  - Gives authority to remove their access
- **APPENDIX**
  - Links to all sub-documents

Site Operational Procedures Policy

- **EGEE-II milestone MSA1.3**
- https://edms.cern.ch/document/726129
- Currently under review, prior to approval as policy
  - **Already approved as a milestone**
- See following slides for full text
  - **The red text is mine (for presentation) and is not part of the policy!**
- Recent concerns/issues discussed and addressed
  - **Relationship with other MoUs and agreements**
  - **Timely deployment of patches**
  - **Where are the project's commitments and guarantees?**

Site Policy (1)

By registering with the Infrastructure as a Site, you and your organization will be deemed to have accepted these operational procedures and policies, unless other agreements on specific issues are in place between the Site and any specific Virtual Organization (VO) or any specific Project, and subject to applicable legislation:

1. *You shall provide and maintain accurate contact information as specified in the Site Registration Policy, including but not limited to at least one Administrative Contact (Site Manager) and one Site Security Contact, in a central repository provided by the Project. Both shall respond to enquiries in a timely fashion, but at least within 3 business days;*

2. *You shall read and abide by the security policies, as published by the Joint Security Policy Group (JSPG) and approved by the Project. You shall periodically self-assess your compliance with these policies, inform the Security Officer of violations encountered in the assessment, and correct such violations forthwith. The Security Officer shall apply appropriate restrictions to the circulation of disclosed information consistent with enforcement and improvement of operational and security policies and procedures.*

Site Policy (2)

3. *Before publishing resource information in resource information systems designated by the Project, you shall ascertain that such resource information is valid and correct to the extent this can be realistically validated. You shall not intentionally publish resource information to resource information systems that is detrimental to the operation of the Infrastructure, or mislead users or their agents into submitting workload, data or information to your Site;*

4. *By accepting workload, data or information from a specific User or VO, you agree to comply with the User or VO requirements as expressed in their respective Acceptable Use Policies, including those relating to accounting and audit data;*

5. *You shall ensure that all the latest security patches and upgrades, and software updates for issues that affect the stability of the infrastructure at large, are promptly applied to all the systems under their control, unless there are strong and justifiable reasons for not doing so;*

Site Policy (3)

6. *Logged information, including information provided to you by Users or by the Project, shall be used for administrative, operational, accounting, monitoring and security purposes only. You should exert due diligence in maintaining the confidentiality of this information;*

7. *Provisioning of resources to the Infrastructure is at your own risk. Software is provided by the Project only as-is, and subject to its own license conditions, and there is no guarantee that any procedure used by the Project is either correct or sufficient for any particular purpose;*

8. *Your Site shall support at least one VO, designated by the Project, for the sole purpose of evaluating the availability of Grid Services at your Site, subject to the provisions made in Article 9. The Project provides to the Site the Acceptable Use Policy and the Security Plan of said VO;*

Site Policy (4)

9. You have the right to regulate and terminate access to Users and VOs at any time for administrative, operational and security purposes. In the case of the Project VO described in Article 8 above, support for the VO must be restored as soon as reasonably possibly. You shall inform the affected Users or VO(s) and comply with the Grid Incident Handling policy regarding the notification of security incidents;

Site Policy (5)

10. *The Project, the Infrastructure management, and their delegates have the right to block your access to the Infrastructure, and to remove or block your resource information from resource information systems, in the case that you consistently fail to comply with this Policy or any of its subordinate Policies (managerial removal), and at any time in case of urgent operational reasons (operational removal). After managerial removal, the mention of your site in both resource information directories as well as in any other publications may be withdrawn. The Project reserves the right to announce, within the Project, any policy violations by your Site, if you fail to respond to and correct such violations in a timely fashion. The Project will facilitate communications between Sites, VOs, Software providers, and Users, in order to enable your Site's compliance with this Policy;*

- This policy shall be signed by an Authorized Signatory of your Organization.

Privacy of User-level Information

Grid Operations and VO managers require…

- **Accounting**
- **Auditing**
- **Logging**
- **Monitoring**

at the VO/group/role/individual user level

Concentrate for now on ACCOUNTING (Usage records)

Legal issues

- EU Directives and national laws relating to…
  - **processing of personal data and the protection of privacy in electronic communications**
- Not allowed to publish data that could be used to track down the activities of a single individual
  - **But can do individual billing (contractual)**
- *personal data* means any information relating to an identified or identifiable natural person
- *Processing* of personal data**: *any operation on personal data, such as collection, storage, retrieval, dissemination* etc…**
- Informed user consent is required
- Issues re data protection, retention, movement (especially across national borders or outside EU)
- User has the right to see their own information
- **Many sites have been unwilling to give access to accounting records at user level**

Grid participant responsibilities

- **Within a site there are usually well defined and well understood processes, rights and responsibilities for network and system admins**
- **Sites (often) restrict access to logs which identify individuals**
- **BUT many Grid participants do not fully understand their responsibilities across the Grid**
  - Users have been identified in mails to lcg-rollout mail list
  - Job details (including DN) are stored and readable
  - These are potentially illegal
- **So the Grid needs well defined policies and procedures**
  - What can and cannot be done and by whom
  - And proper training of sys admins, GOC staff and VO managers

User consent

Grid AUP says…(accepted during registration with VO)

- ***Logged information, including information provided by you for registration purposes, shall be used for administrative, operational, accounting, monitoring and security purposes only. This information may be disclosed to other organizations anywhere in the world for these purposes. Although efforts are made to maintain confidentiality, no guarantees are given***

- **So the User has given informed consent**

- **Together with a policy document on personal data management, should be enough to convince sites to allow access to the appropriate logs**

Accounting Requirements

WLCG requires (OB)

- **All Grid sites providing resources to WLCG**
- **to provide accounting information for all jobs run by members of the LHC VO's**

The LHC VO managers require

- **Access to accounting at individual user level**
  - i.e. X.509 DN included
- **Again from all sites providing resources**
- **Timely accounting (for monitoring) useful?**

Assume other VO's have similar requirements

- **We need a technical solution and agreed policy that allows all this to happen legally**

Accounting Data Policy

- **Several presentations to LCG GDB on this topic**
  - Most recently by John Gordon in June 2006
- **We need feedback (from VOs) on the proposal for treatment of user-level accounting**
- **Aim to be technology independent**
  - APEL and DGAS
- *Once all agreed we can finalise the policy document*
- **One new requirement has recently arisen**
  - UK GridPP Oversight Committee
    - Requests to see details of how many different users are using the GridPP resources (UK and not-UK)
- ***Is there a general requirement for the Operational infrastructure to see user-level accounting?***
  - In the past we have always assumed just VO's
  - Feedback welcome!

Draft Accounting Policy

- **All EGEE sites are required to provide to the GOC accounting DB an accounting record per job in the agreed format (~GGF schema)**
  - Today via APEL
  - Policy must also apply to DGAS
- **This record not only specifies the VO but includes the User DN (encrypted)**
- **These records must be transferred to the central GOC accounting DB on a regular basis**
  - How frequent? (every day?)
- **VOs are not allowed to do their own accounting or bookkeeping (except with agreed policy)**
  - Same legal problems of monitoring individual user activity

Accounting Policy (2)

What is stored and where?

- Grid service/batch system log files stored locally
  - These include unencrypted user DNs
  - Audit requirements policy requires these to be kept
- Regular processing -> accounting record per job (locally) in the agreed format
  - In APEL, stored on the site R-GMA MON box
- User DN field must be encrypted
  - use one standard accounting public key
  - Encrypt “DN + random string” to avoid analysis of the encrypted DN
- Records regularly -> RAL GOC accounting R-GMA MON box
- Subsequently -> GOC accounting database

Accounting Policy (3)

- **USER DN will be decrypted during this transfer (to GOC DB)**
- **Individual accounting records are aggregated into summary records in GOC accounting database**
  - For example…
  - one record/VO/month/site
  - one record/VO/month/user
  - Or should aggregation be per week?

Accounting Policy (4)

Access rights - Who can read?

- **R-GMA today has Authentication but not Authorization**
  - Anyone with a valid certificate from an IGTF approved CA can read the RAL GOC R-GMA data
  - User DN must remain encrypted in R-GMA
- **The decryption key is only known by authorised GOC staff**
- **The GOC accounting Database (including decrypted user DNs) is only readable by the authorised GOC staff and (via a web interface) by authorised VO staff**
- **General read access (web front end) to the accounting data must not include user-level info**
- **Any user has the right to access his/her own accounting records**
  - mechanism to be defined
  - By request to GOC or by authenticated web access?

Accounting Policy (5)

Access for what purposes?

- To aggregated accounting data (no user data)
  - **For all bona-fide accounting purposes**
- To aggregated user-level data
  - **Only for operational and monitoring purposes**
  - **Must preserve the confidentiality**
    - **No web publication**
    - **No discussion on public mail lists**
    - **Information must remain within authorised management of GOC and VOs**

Accounting Policy (6)

Retention period?

- **The detailed records**
  - one per job including encrypted user DN
  - Both in R-GMA and GOC accounting DB
  - Will only be retained for as long as needed
  - Propose deletion after 12 months
    - annual resource planning cycle
- **Aggregated summary data**
  - With no user DN – no deletion defined
  - Summary per user – propose deletion after 12 months

Accounting Policy (7)

Publication of accounting data?

- **Access via web front end(s)**
- **No access to individual job records**
- **Aggregated summary data (per VO/site/month)**
  - World readable (or need a certificate?)
- **Aggregated summary data (per VO/user/month)**
  - Not world readable
  - Authorised GOC and VO managers may read
    - Must not publish anywhere else
    - Must not store analyses anywhere else

Accounting Policy (8)

How is the data protected from illegal or accidental disclosure?

- **R-GMA data needs no special protection**
  - Needs a valid IGTF Grid certificate – Authenticated access
  - Most fields not private
  - User DN is encrypted to prevent disclosure
- **GOC Accounting DB (MySQL)**
  - Authorised access only
  - Protected via normal MySQL and OS/File system access control
  - But no guarantees, of course

Site feedback on legal issues

- 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.**
- **Accounting requirements.** The policy must contain a list of strong arguments for the necessity of collecting accounting data and the purpose thereof

Feedback (2)

- **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).
- **Routine usage of the data.** The policy should describe the routine usage of the data as completely as possible.
- **Misuse of and suspicion on wrong accounting data.** The policy should describe a procedure to be followed in such cases.

Who needs to see user level records?

- The **VO** is the group of people who are members
- **Members** are added or deleted to/from the VO. The **VO Manager** 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
  - **only person who has access to all data in the accounting database (for the VO) including the DN.**
  - **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.**

Next steps

- **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'**
- **Work on the policy will continue …**
  - We are awaiting feedback from VOs
- **Anyone holding user level information should be interested in this.**
  - including all the various monitoring and logging systems

JSPG future plans

- **EGEE SA1 milestone (MSA1.7)**
  - *Update Security Policy* (Month 8 = Nov06)
- **VO AUP**
  - OSG working on this
  - I believe EGEE should adopt this approach
    - Requires VO's to accept responsibility.. Will this work?
- **Minimum requirements for VO membership services**
  - JSPG agreed this would be very useful (work not started)
  - Work with OSG on general Service AUP
- **Risk Analysis and Security Plan**
  - OSG work presented here at EGEE'06
  - EGEE long term strategy to do something similar
    - do we have resources to work with OSG on this?
- **WLCG security emergency plans**
  - Define communication & decision processes (started)

JSPG Meetings, Web etc

- **Meetings - Agenda, presentations, minutes etc**

http://agenda.cern.ch/displayLevel.php?fid=68

- **JSPG Web site**

http://proj-lcg-security.web.cern.ch/

- **Membership of the JSPG mail list is closed, BUT**
  - Requests to join stating reasons to D Kelsey
  - Volunteers to work with us are always welcome!
- **Policy documents at**

http://cern.ch/proj-lcg-security/documents.html

Grid AUP (1)

By registering with the Virtual Organization (the "VO") as a GRID user you shall be deemed to accept these conditions of use:

1. ***You shall only use the GRID to perform work, or transmit or store data consistent with the stated goals and policies of the VO of which you are a member and in compliance with these conditions of use.***
2. ***You shall not use the GRID for any unlawful purpose and not (attempt to) breach or circumvent any GRID administrative or security controls. You shall respect copyright and confidentiality agreements and protect your GRID credentials (e.g. private keys, passwords), sensitive data and files.***
3. ***You shall immediately report any known or suspected security breach or misuse of the GRID or GRID credentials to the incident reporting locations specified by the VO and to the relevant credential issuing authorities.***

Grid AUP (2)

4. *Use of the GRID is at your own risk. There is no guarantee that the GRID will be available at any time or that it will suit any purpose.*
5. *Logged information, including information provided by you for registration purposes, shall be used for administrative, operational, accounting, monitoring and security purposes only. This information may be disclosed to other organizations anywhere in the world for these purposes. Although efforts are made to maintain confidentiality, no guarantees are given.*
6. *The Resource Providers, the VOs and the GRID operators are entitled to regulate and terminate access for administrative, operational and security purposes and you shall immediately comply with their instructions.*
7. *You are liable for the consequences of any violation by you of these conditions of use.*

Example VO AUP

This acceptable Use Policy applies to all members of <The VO> Virtual Organization, hereafter referred to as the VO, with reference to use of the LCG/EGEE Grid infrastructure, hereafter referred to as the Grid. The Geant4-Spokesman, <name> owns and gives authority to this policy. The goal of the VO is to validate the software they provide to their users (HEP experiments as ATLAS, CMS, LHCb, Babar, etc, Astrophysics applications, biomedical communities) twice per year within the Grid environment. This procedure should cover a wide range of parameters and physical models which are high CPU demanding. At the same time they are planning to use regularly the LCG/EGEE resources to make analysis and studies of their toolkit. Members and Managers of the VO agree to be bound by the Grid Acceptable Use Policy, VO Security Policy and other relevant Grid Policies, and to use the Grid only in the furtherance of the stated of the VO.