Study on Authorization

Christoph Witzig, SWITCH
(witzig@switch.ch)

JRA1 Oct 25, 2007

www.eu-egee.org

Outline

- **Introduction**
  - Goal of this study
  - Priorities of the study
- **Requirements from the experiments and sites**
- **First ideas --> discussion**

  Note: These are ideas and I am now hoping for feedback (in this meeting or by email afterwards)

Goal of the Study

- **Task by C.Grandi to look into authorization (authZ) in gLite with the goal to specify design for "authorization service" work item in EGEE-II/-III**
  - EGEE-III proposal: authZ service: Nikhef, UvA, CNAF, SWITCH
- **Should specify work in 2008 / early 2009**
  - Comment: should be fully deployed within lifetime of EGEE-III
- **Deliverable is a proposal with clear recommendations based on input of many people (experiments, SAx, JRA1) to be accepted/rejected by TCG**

Timeline

- September / early October: requirement gathering

  PLEASE let me know to whom I should talk to

- mid-October - late Nov: working out the recommendations and a proposal of the design

- Discussion at MWSG meeting in December

- Presentation and decision in TCG in January

Priorities

List of priorities in order (as approved by TCG):

1. **Should fix some of the limitations of the current authZ framework**

2. **Introduce new features to the extend that they are needed by the**
    1. Experiments / VOs
    2. Sites / SAx
    3. JRA1

3. **Interoperability**

4. **Use of standards if possible**

Disclaimer

In the following I will present some first initial ideas to be discussed

Some of them may end up being recommendations, some may not

Feedback from developers most welcome

Outline

- Introduction
  - Goal of this study
  - Priorities of the study
- **Requirements from the experiments and sites**
- Review of existing authorization mechanisms
- First ideas --> discussion

Requirements from VOs

- **So far spoke with CERN experiments**
  - There are other VOs than LHC experiments

- **Personally I believe many new requirements from VOs will appear once physics data is available**

Requirements from VOs (2)

- **They want a simple system that works and gives them most of the control of the system**

- **< 50 groups, a few roles**

- **Roles:**
  - production: to run production jobs
  - lcgadmin: to install software
  - soft-valid: to run monitor jobs

- **Groups:**
  - Currently very few groups
  - I would expect this will change once we have production data

Requirements from VOs (3)

- **Differerent kind of jobs:**
  - Read data from SE, process and write data to SE
  - Need access privileges for software installation at sites
  - Manipulate data in SE and catalogues which may require special priviledges

Requirements from Sites

- **They want a simple system that gives them <u>all</u> of the control of their site**
- **From their perspective they want minimal installation and maintenance while supporting VOs**
- **They want**
  - simple, clear rules for installation and maintenance
  - configuration files that are simple to set up, change rarely and are intuitive to understand
  - understand and control user mapping
  - have a grasp what the user jobs are doing (security wise)
- **We should be aware that the site administrators community is rather diverse**
  - Knowledge in grid ranges from expert to novice

Current authZ Mechanism

- **authZ = permission to access a resource based on a set of attributes**

- **Basic mechanism in gLite:**
  - Proxy certs with VOMS extensions
    - DN, pFQAN, sFQANs
      1. *identity of the user*
      2. *membership in VO (and its subgroups)*
      3. *role (dynamically chosen by the user)*
  - Use of this information by different algorithms at different places in the middleware

Outline

- **Introduction**
  - Goal of this study
  - Priorities of the study
- **Requirements from the experiments and sites**
- **First ideas --> discussion**

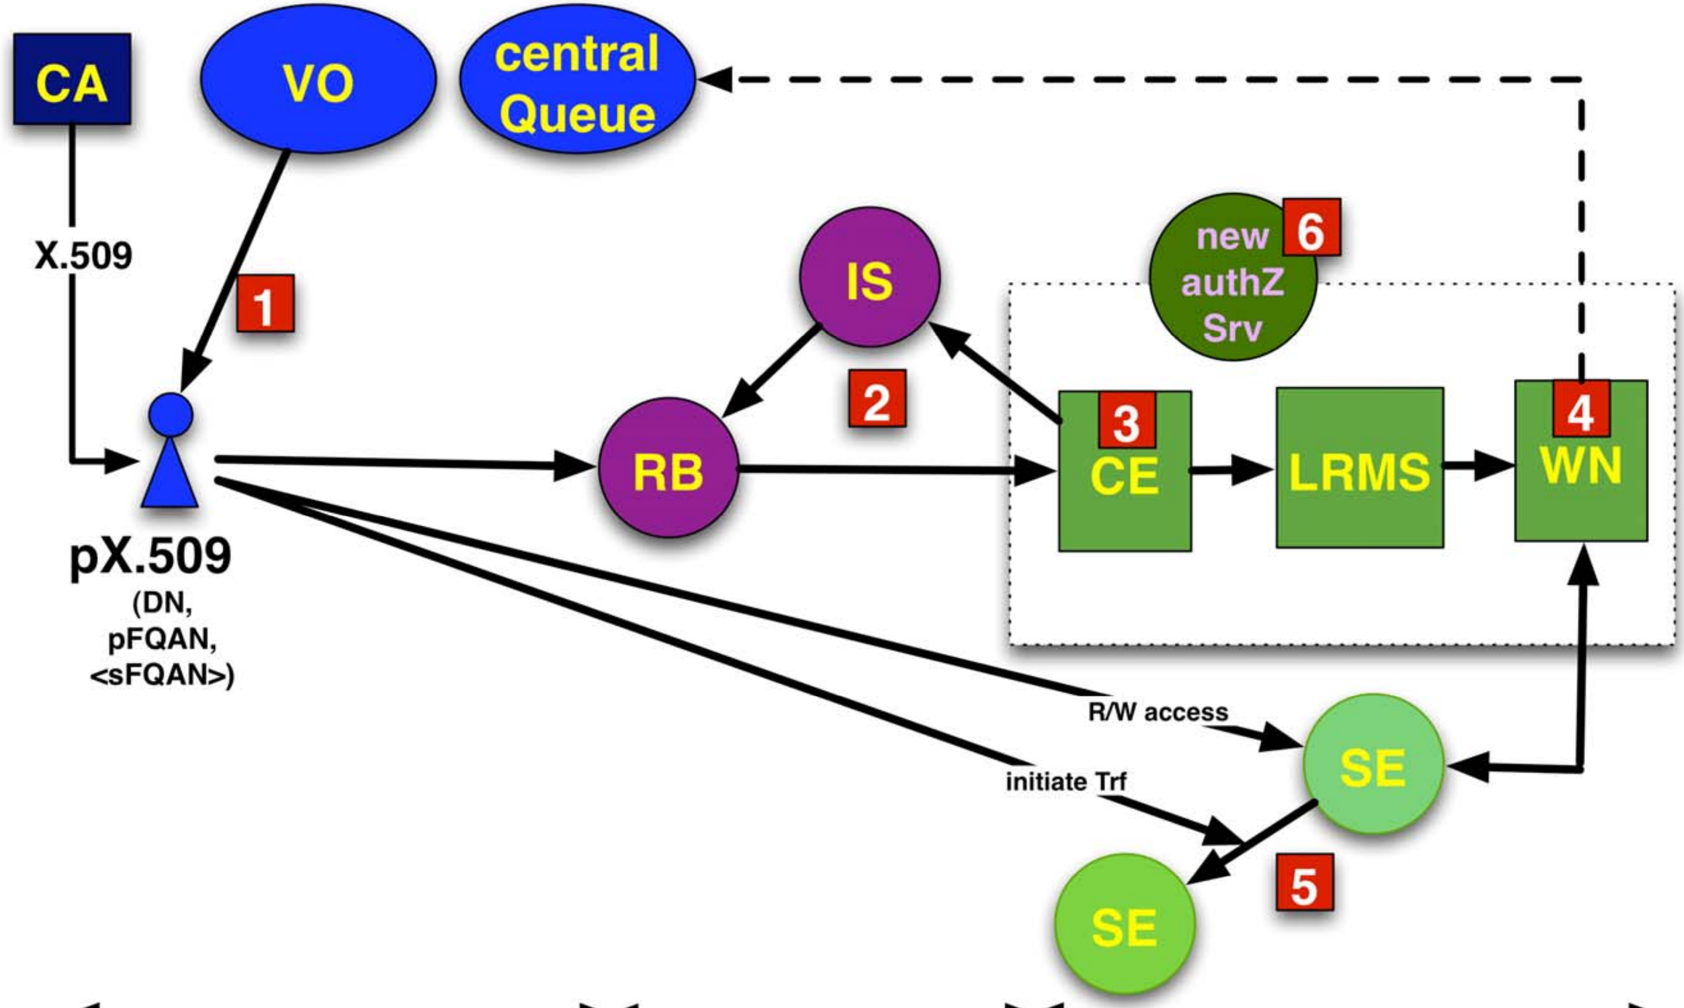

🟦 VO

Req: changing with time

🟪 Central Services

🟩 Sites

Req: local site control

List of Topics

- Documentation
- Pattern Matching Rules for FQAN
- VO/User Level
- Inside the CE
- Resource Broker and BDII
- Data Management
- WN and new authZ service

- Ideas are clearly marked in the slides and are numbered

Documentation

- **Recommendation:**
  - Have a document which describes the authorization mechanisms
  - Comment: will be one of the deliverables of this study

Pattern Matching Rules

- **Source of many confusions**
  - Example recent GDB discussions
- **There is now a document describing the rules and a set of matching functions (in C and Java)**

- **Idea #1: Have a standard library, which**
  - Developers must use when matching FAQN patterns OR
  - Existing code is reviewed and tested against a set of patterns (where replacement with library functions is not feasible)

- **Idea #2: Change rules for wildchar**
  - Wildchar are no longer supported OR
  - Wildchar are only allowed after the trailing “/”
    - E.g. /atlas/analysis/* and not /atlas/*/higgs or /atlas/ana*

1. VO/User Level

- **VO wants to add "attributes" to the user**
  - Is VOMS groups/roles enough?
  - What kind of information does it want to pass along to the user?
  - What kind of information shall be user specify at submission/access time?

- **Idea #3: Support only groups (no more roles)**
  - Motivation:
    - We anyway match only on strings so the effect of having a group and a role is equivalent of just having a (unique) string
      - *E.g. /atlas/role=production is equivalent to /atlas/production*
    - FQAN matching becomes easier to understand

- **Idea #4: Give the user more control over FQANs**
  - Currently only the primary FQAN counts
    - And user may not be aware of it
    - New proxy is needed for new primary FQAN
  - Consider FQANs as a set of keys (that are anyway always present as a complete set in the AC of the proxy) from which the user chooses which one to take (see also idea #12)

3. CE

- **Sites want to authorize the user based on a set of attributes**
  - DN/FQAN --> uid/gid(s) --> share LRMS
    - Links authZ info with scheduling
  - Site administrators want to
    - retain complete local control
    - Clearly understand the mapping to uid/gids
    - Simple management
      - *Consider >1 CE per site*
  - VOs want “intelligent” scheduling at the site
    - Mapping of FQANs sometimes statically, sometimes dynamically

Inside the CE

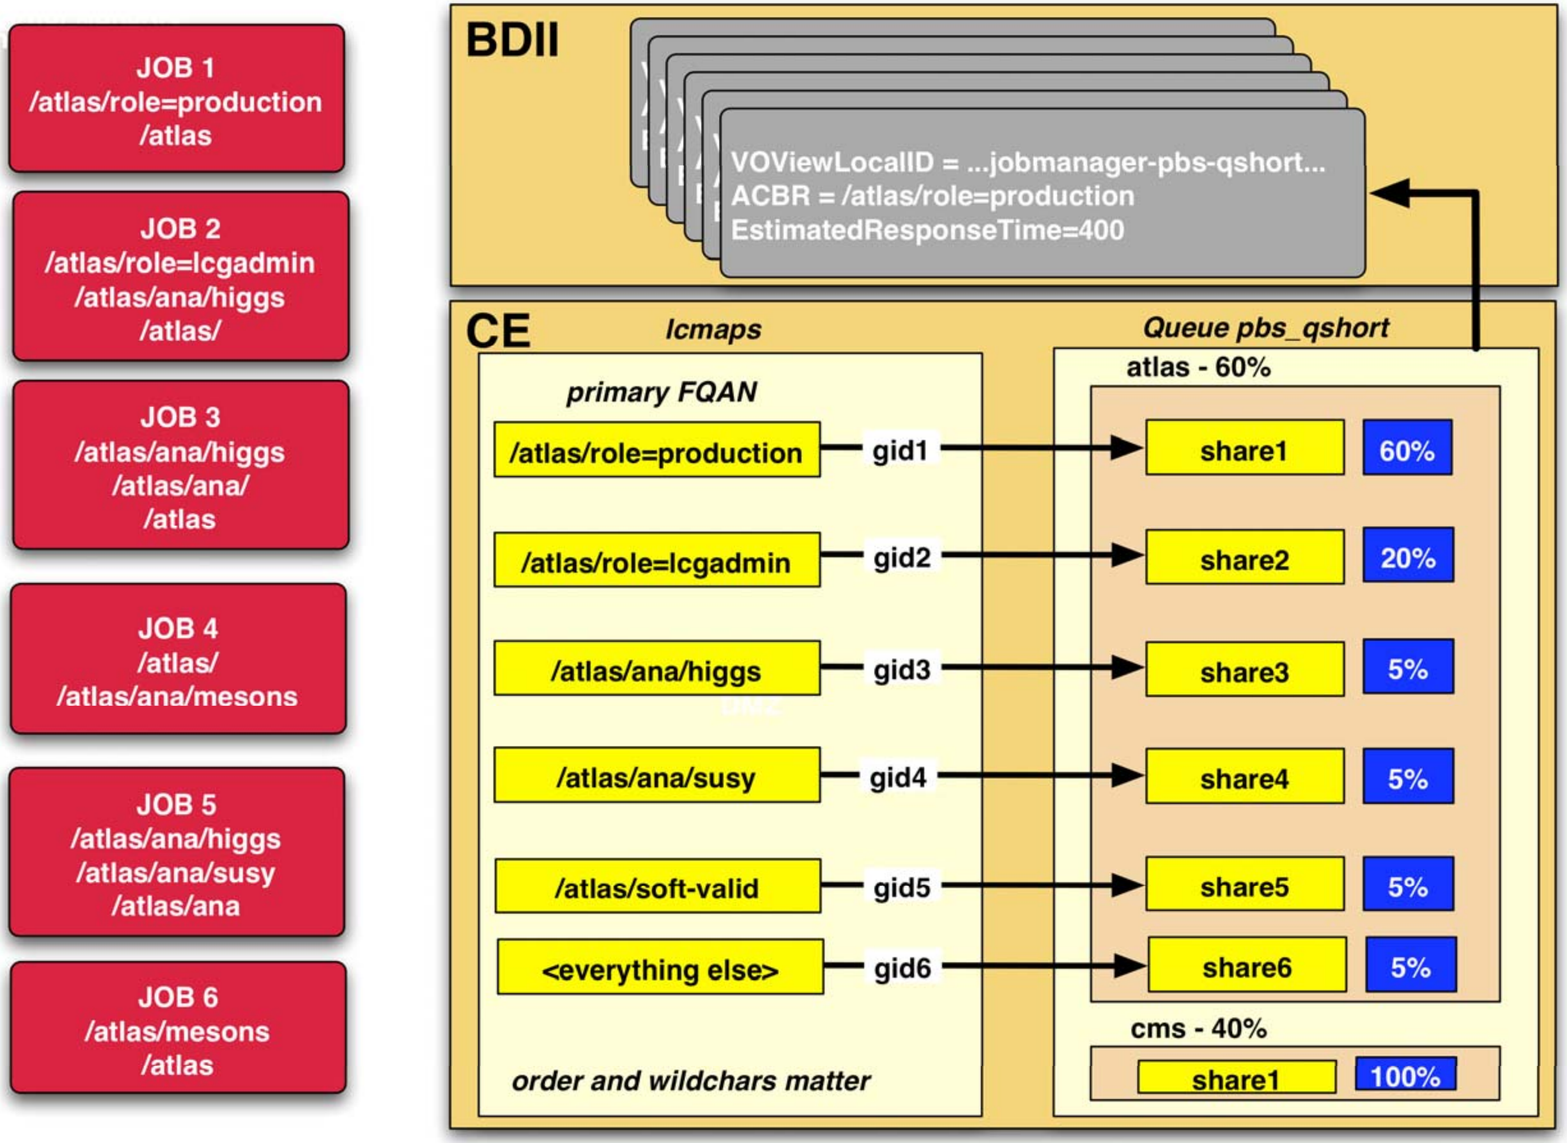

JOB 1
/atlas/role=production
/atlas

JOB 2
/atlas/role=lcgadmin
/atlas/ana/higgs
/atlas/

JOB 3
/atlas/ana/higgs
/atlas/ana/
/atlas

JOB 4
/atlas/
/atlas/ana/mesons

JOB 5
/atlas/ana/higgs
/atlas/ana/susy
/atlas/ana

JOB 6
/atlas/mesons
/atlas

BDII

VOViewLocalID = ...jobmanager-pbs-qshort...
ACBR = /atlas/role=production
EstimatedResponseTime=400

CE

lcmaps

primary FQAN

/atlas/role=production — gid1 → share1 60%

/atlas/role=lcgadmin — gid2 → share2 20%

/atlas/ana/higgs — gid3 → share3 5%

/atlas/ana/susy — gid4 → share4 5%

/atlas/soft-valid — gid5 → share5 5%

<everything else> — gid6 → share6 5%

order and wildchars matter

Queue pbs_qshort

atlas - 60%

cms - 40%

share1 100%

Inside the CE (2)

- **LCMAPS currently takes the first match**

- **Idea #5: should it take the most specific match?**
  - Regardless of location within the mapping file?
  - How do you define “most specific match”?

Inside the CE (3)

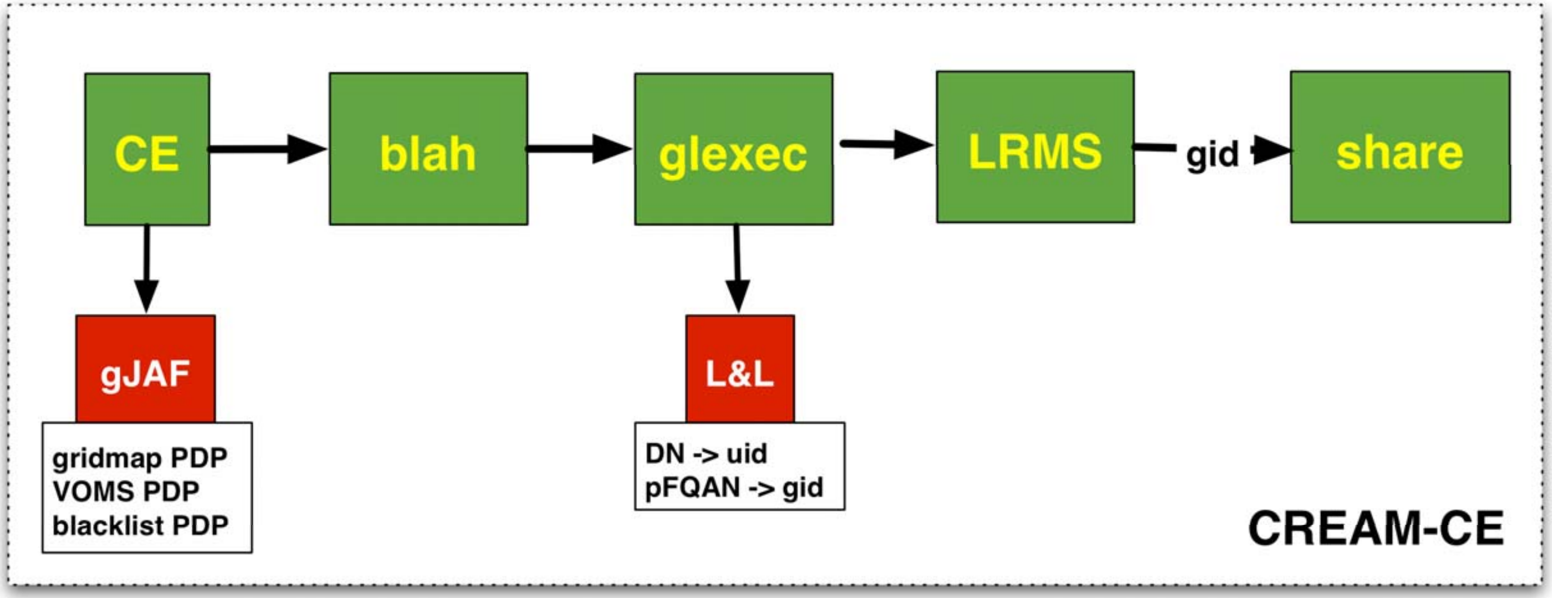

- **Comment: Does it make sense to have two authZ frameworks? (gJAF and L&L)**
  - Architecturally?
  - Further development work and code maintenance?

- **Idea #6: Drop one or unify them (e.g. through authZ service)**

authZ and Scheduling (1)

- **Idea #7: remove coupling of authZ data and scheduling at CE**
  - Two implementation ideas: next slide
  - Comment: Don’t make pilot jobs this obsolete? Yes - but
    - Not all VOs will have their own scheduling system
    - Not all sites may support pilot jobs
    - We don’t know (yet) the security implications of pilot jobs

authZ and Scheduling (2)

- **Two ideas:**
  - Preset hierarchical model of fixed shares (idea #8)
  - Dynamic assignment of shares (idea #9 - next slide)
  - Both are ideas -> **further work needed to clarify implications**

- **Idea #8:**
  - Have a (small) predefined set of shares within a VO that never change (e.g. gold, silver, bronze, lead, carbon, hydrogen, ….)
  - At the level of the VO assign attributes to groups that reflect these “QoS”

- **Comments:**
  - Cannot easily change the predefined shares (static solution)
  - This is independent of the user mapping (site requirement)
  - Matchmaking must take supported VO information and VO view of predefined shares into account

authZ and Scheduling (3)

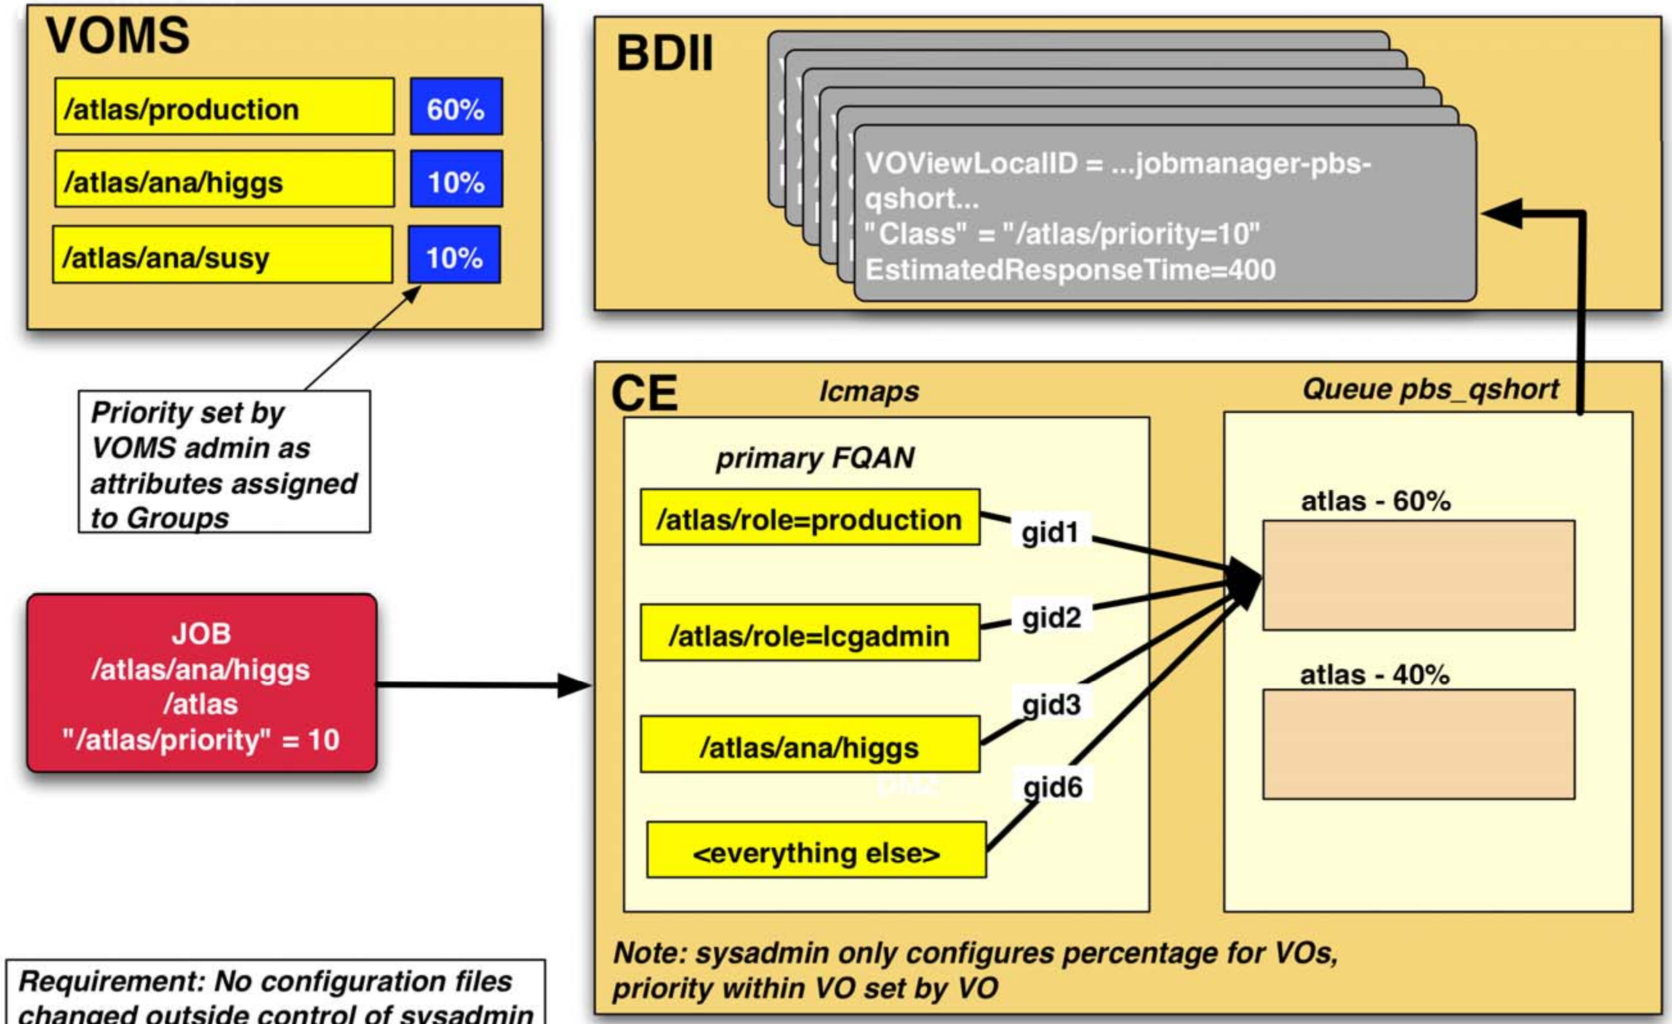

Priority set by VOMS admin as attributes assigned to Groups

Note: sysadmin only configures percentage for VOs, priority within VO set by VO

Requirement: No configuration files changed outside control of sysadmin

Question: can this be done with supported LRMS?

2. RB and IS

- **Resource Broker**
  - Needs info from sites for “push” model
    - IS is mainly for service discovering, has limited capabilities for giving complete overview of situation at CE
    - How should situation at CE be published?
    - Pull model wanted - which problems does that solve/raise?
  - Should not send job to CE where it will be rejected (must be “site policy aware”)

- **Consistent policy between RB and CE is a “must” requirement**

- **How? - Ideas**
  - Idea #10: VO views must guarantee full view of all mappings OR
  - Idea #11: RB must be able
    - (somehow) to know (and incorporate into match-making) what will happen at the CE through
      - *New service?*
      - *Extracting the map file?*
    - OR tell the CE which FQAN to take and know that the CE will do that

5. DM (1)

- **Access rules (security) models differ in DM**
- **I only considered DPM**
- **Assumptions that Castor and dCache will follow DPM model**
  - Is this correct?

- **Idea #12: Use different FQAN for DM than for job submission**
- **Extension: Use different FQAN for different operation on the Grid**
  - Example: FQAN for job submission, FQAN for DM, FQAN for Catalogue access
  - Determined by user at job submission time

- **Analogy: FQANs in AC are a set of keys of which one or the other is selected for a given operation on the grid**

- **Note: consequence is that the primary FQAN is no longer significant if FQANs are selected for operation (could be default value if no FQAN is selected)**

5. DM (2)

- **Motivation:**
  - DM has different requirements than jobs (data resides much longer on disk than job execution)
  - Storage is expensive - may need to assign different lifetimes to files written by different users (example files written by atlas user of a given institution vs atlas user of other institutions)

- **Requires:**
  - Support in JDL file
  - Support at command line argument (from job script in bash, python, …)

5. DM (3)

- **Idea #13: Assign different quotas to different FQANs in a hierarchical way**
- **Comment: Is this really a authZ question?**
- **Comment: There was also a requirement on admin authZ at FTS (which I haven’t followed up yet)**

4./6. WN and new authZ Service(1)

- Requirement of glexec on WN
- Web service based authorization service
- Document of Yuri

4./6. WN and new authZ Service(2)

- **Concerns (purely personal):**

1. Performance
   - We must be aware of the possibility that poor performance may lead to the situation that it will not be used by sites
2. What are really the goals here?
   - Interoperable with what and whom?
3. Choice of SAML library?
4. Have to agree
   - on interface specification that allow different implementations
   - within EGEE how this relates to current and future code, in particular SCAS service and G-PBox
5. I believe it's wrong to say that development of this service is so urgent that concern 1.-4. cannot be addressed properly now (i.e. before the end of the year)

   unless one agrees that this is a short term solution that might undergo drastic revisions

Design on authZ Service

- **New Authorization Service (from C.Grandi's slides)**
  - Uniform authorization and policy management in gLite
  - Compatible with SAML and XACML standards
  - Built on the experience of previous systems
    - LCAS/LCMAPS, SCAS, G-PBox, gJAF
  - Full design to be completed by Dec. 2007, implementation will start in 2008 and continue in EGEE-III
  - Not constrained to the use of any existing implementation
    - though recommended for the sake of economy

G-PBox (1)

- **Promises to synchronization RB - CE**
- **Uses XACML**
  - Promise of more complex policies in consistent way
  - “hard to read” -> acceptance by sites?
- **Policy change**
  - From VO -> CE
  - CE -> VO
  - What mechanism is really needed here? G-PBox may or may not be the right solution

G-PBox (2)

- **Comment: I have currently no opinion on G-PBox and we will have to look at it more closely**

Summary

- **Presented list of ideas (1...13)**

- **Feedback on**
  - these ideas or
  - new ideas

  welcome at any time by email (witzig@switch.ch)

- **Next steps:**
  - In November:
    - Write draft version of authorization document
    - Formulate recommendations based on feedback on ideas presented
  - In December:
    - Expect feedback on recommendations
    - Discussion with OSG at MWSG
  - In January:
    - Final version of the authorization document
    - Present list of recommendations to TCG