



Support for VOMS roles in production services

WLCG Workshop, CERN

23 Jan 2007


Andrea Sciabà (CERN) - Experiments

Jeff Templon (NIKHEF) – Sites


Chris Brew (RAL) – Tier 2

Maarten Litmaath (CERN) - Deployment

(convenor Dave Kelsey - RAL)



Andrea Sciabà CERN



Usages of VOMS for the experiments

- To implement a fine-grained management of permissions and privileges in the Grid
 - Access to computing resources
 - Job priorities!
 - Access to data and storage resources
 - Management of VO services
 - Accounting




VOMS and job priorities

- **Basic need to define how many CPU resources should be assigned to different activities**
 - e.g. give 60% of CPU to production at site X for the next three months
- **Different possible prioritization mechanisms**
 - Global VO task queue: the VO decides which jobs should be submitted first; easier for sites
 - Local prioritization: jobs are submitted equal, the priority is set by the local batch system based on the primary VOMS FQAN in the proxy; site configuration involved
- **How many VOMS groups/roles are expected for jobs?**
 - In 2007: not much more than distinguishing "analysis" and "production"
 - 2008 -: tens of them
- **How often should the priorities change?**
 - In 2007: maybe once/month, easy manageable even if done by hand
 - 2008 - : using some automatic tool to propagate the information to sites
- **What kind of support for VOMS and job priorities should the middleware provide?**
 - Publication to the information system of VOMS-specific info (VOViews for VOMS FQAN, not just for VOs)
 - The WMS should use the above information for match-making and ranking
 - **But it is possible to live without this in the short term**



VOMS and data management

- VOMS-aware storage and catalogues
 - Specify quotas for different VOMS groups
 - Associate special roles to write access
- VOMS-aware data transfer tools
 - prioritize data transfer jobs based on VOMS group/role
 - e.g. allocate 70% of T2→T1 bandwidth to MC transfers
- What is used now?
 - VOMS access control lists in LFC and DPM
- What should be available in the future?
 - Short term: actual permissions should take into account all the groups the user is in, not only the primary group in the proxy
 - Long term: VOMS support in all SRM storage systems




VOMS and other services

- **Software installation**
 - Special role for installing VO software at sites: first real use of VOMS
- **VO-specific services**
 - Roles for service administrators to change service configuration, etc. opposed to "normal user" access
- **Accounting**
 - Build resource usage statistics as a function of the VOMS FQAN (not only of the VO)
 - e.g. how many resources were used by production jobs at site X?



VOMS group tree

- Many possible group hierarchies. Some examples:
 - Regional groups
 - `/vo/regionX/siteY`
 - Activity-related groups
 - `/vo/analysis/StandardModel`
 - `/vo/reconstruction`
- It is expected to have in future $O(100)$ distinct group/role combinations



VOMS : site perspective


Jeff Templon

NIKHEF

wLCG workshop


23 january 2007





Where are we

- VOMS works for SW installation ...
Longstanding issue
 - One was either an SGM, or a user, but not both ... Solved
- Elementary separation of storage (all or nothing is all that is guaranteed, like old SGM situation)
- Website access etc. via certificates




Where are we almost

- **Basic job priorities**
 - Separation of shares via groups / roles
 - Publishing of ERTs per group / role
 - Matchmaking due to group / role using this info




Where we ain't ...

- **Accounting via group role**
 - How do exp'ts and sites know it's working?
- **Information management**
 - Scores of combinations group/role -> IS explosion??
- **More flexible DM VOMSification**
 - Better than all or nothing, ie a production manager could *also* write into “normal” ATLAS share ...



And will it be consistent??

- Lots of decisions to make ...
 - How to choose “most appropriate” group / role in case multiple matches?
 - How to make sure that all subsystems make same choice?
 - How to limit the damage in the inevitable case that they don't?
- Make sure that mapping VOMS=> unix gids/uids does not create “hidden” limitations



Chris Brew
RAL




A Tier 2's Concerns

- Proposal for VOMS based Scheduling (as I understand it)
 - VOMS Groups → Unix Groups → MAUI
- On CE:
 - Proliferation of Pool Accounts and Groups
 - Already 1000's of pool accounts. How many will I need how often will I need to create new ones?
 - Frequency of updates to maui.cfg
 - How often will I have to be changing scheduling priorities?
 - Maintainability
 - How can I tell if It's working?
- On SE:
 - Will I have to set up separate pools/endpoints for each Group?




Maarten Litmaath CERN




VOMS & Deployment

- schedules must be driven by EMT and TCG !
 - experiments have to ensure the right issues are on the agenda
- new YAIM on CTB allows for special cases other than sgm and prd
 - allows new groups and roles to be mapped differently
 - none has been added so far




VOMS & Jobs

- Job Priorities WG batch system recipe
 - where is it? ([link to document](#), comments welcome, room for DM stuff & LSF included)
 - is the TCG document of July 20 the latest?
 - needs new Torque + Maui, in certification
 - needs new lcg-info-dynamic-scheduler rpms, in certification
 - needs changes in YAIM, not done yet at least for LSF



VOMS & Accounting

- accounting for groups and roles almost ready
 - LCG-CE patch awaiting YAIM update
 - APEL patch expected in a few days



VOMS & Data Management

- supported by LFC and DPM
 - support for multiple groups and roles expected in a few weeks
- supported by dCache to some extent, via gPlazma callout
 - LCMAPS plugin to be written, not clear by whom
- not yet supported by CASTOR, still using grid-mapfile
- not seen as a priority for this year
 - separation of production managers from users is working
- FTS 2.0 supports VOMS, to be released in a few weeks



Discussion