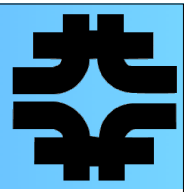

US-CMS Privilege Management

April 23, 2003

Preliminary Privilege Ideas

US-CMS facility requirements for privilege management and authorization are different from other facilities at FNAL

- There will be many users but most will have never been to FNAL and many will have no association with us
  - CMS is huge.
- A lot of input will come from the VO about who has is a member, what the priorities are, what sub-groups are authorized to perform certain tasks
  - At the same time the facility needs to protect itself

Most of what we're thinking about today pertains to "Grid" users. The expected local user requirements are within the capabilities of existing techniques

The primary areas for CMS are

- Authorization and privilege that enables the user and the VO to work
- Authorization and privilege that allows the facility to function securely

Two models for Privilege Enforcement

There are currently two models of where to handle the privilege and authorization enforcement

- The VO model
  - In the VO model authorization decisions are handled centrally at the VO and work is doled out computing facilities.
    - The VO can be entirely trusted.
- The Facilities/User Model
  - The policies still come from the VO but the enforcement is handled by facility services
    - A user asks for access rights to a dataset in his/her role as physics coordinator, the dataset is registered as belonging to the coordinator and access is granted by the site
    - The user cannot be trusted and must be checked

The will probably be elements of both models in the system that CMS will eventually deploy

VOs Activities

VOs need to be able to establish priorities for VO activities

- A VO needs a centralized work managers that assigns the priorities
- Or Users need to be able to specify intentions when they arrive at a site. The VO needs to be able to control the priority
  - "I am a performing re-reconstruction in my role as Higgs coordinator"
  - VO defined policy is then enforced on the site

At the moment US-CMS has neither. Incoming users are mapped to group accounts for the VO, we have no idea what their applications are or how they should be prioritized with respect to other VO members

We currently have the ability to identify which user ran a specific process. We don't have enough information to provide finer grained auditing and accounting

In the long term CMS needs to establish how VO authorization is handled. In the short term we interested in using extended proxies

Data Access

VO and users need the ability to control who has access to data (files, databases, storage)

- Central VO controlled files and areas are easier. They are likely read-only except by people with defined roles.
- User data is much more difficult and one of the places CMS would like to work
  - Access rights will change, it needs to be audited and quota enforced, it is transient

With the current account mapping it is difficult to tell who made the file and who should have access to it.

- Even finer grained role based mapping is probably not sufficient
- Need to arrive at a level where a file is identified with an individual
  - Ideally access would also be controlled at the level of the individual
    - I user X allow you user Y access to my dataset for a period of time

Facilities

Whether the request comes from the VO or the user directly, the facility needs to be protected against abuse.

- The facility needs to be able to identify and stop users and activities
  - Or the facility needs to get in a mode where we aren't responsible for activities, which seems unrealistic
    - Modules like SAZ are likely to remain relevant
- CMS is interested in investigating authorizing specific services based on user, role, and process
  - Network services are potentially the most dangerous so they would be a good place to start
  - Other services to be investigated