NeuroLOG ANR-06-TLOG-024
Software technologies for integration of process, data and knowledge in medical imaging

Financé par ANR

NeuroLOG Security Policy proposal

Technical document

Authors: J. Montagnat (I3S)
D. Godard (Visioscopie)

Summary: This document describes the security policy proposed to control the access to sensitive resources within the NeuroLOG project. It also drafts a possible implementation of this policy.

Document layout

1. Summary of proposed security

The policy described in the next section aims at fulfilling the security requirements of the NeuroLOG project for data access control. The proposed solution also intends to be as lightweight and easy to deploy as possible. To summarize, the main aspects of the proposed policy are:

1. The security policy is administrated locally on each site by one site administrator. There is no global administrator of the distributed platform.
2. All users are securely identified and registered into the system. The local site administrator
3. Data access is controlled at the level of groups of users. At least one group will be created for each site (by default all data registered into one site will be accessible to the site group only). As many additional groups may be created as needed. A group is created and owned by one site but known from the other sites: there is a single administrator for one group (the group local site administrator) but users from other sites may be registered into any group. A site administrator may decide which data is accessible to which group.
4. Accesses to local data will be traced individually in a non repudiable manner on each site.

This policy ensures that each site controls its data: local data access is under the responsibility of the site administrator. However, groups containing users from different sites allow data sharing among different sites.

2. Security policy description

Sensitive medical data will be deployed on different sites over the distributed NeuroLOG platform. The NeuroLOG Security Policy (NSP) aims at fulfilling two *a priori* antagonist roles:

- To make data exchanges among users from different site possible;
- To ensure that each site solely controls the access to the data it owns.

To implement the NSP, all system users are authenticated through non-repudiable nominative certificates. Each user is registered into one site (and thus known by the system) by the site administrator. Access to data is controlled at

the group level: as many groups as needed may be created and data files are individually controlled by group.

Complementary to the NSP, all data exported from a site will be anonymized and encrypted prior to transmission for protection as detailed in NeuroLOG deliverable 3.

2.1. Users, administrators and groups

Each system operator, whether normal a user or an administrator, is identified by her/his individual X509 certificate. Such a certificate contains a non repudiable Distinguished Name (DN) similar to:

```
/O=GRID-FR/C=FR/O=CNRS/OU=I3S/CN=Johan Montagnat
```

This certificate thus identifies the user name and institution. Several credentials will be used in the NeuroLOG platform: login/password identification, CPS (Carte Professionnelle de Santé), SQL92 identifiers and grid certificates. The system will ensure the proper mapping of a single user DN to all these credentials to interoperate with the services.

On each site, a single user gets the administrator privileges allowing her/him to:

1. Register or unregister other users into the local site. Only users with a valid certificate belonging to the local site may be registered.
2. Change the administrator privilege recipient.
3. Create groups and populate groups with user DNs.
4. Grant group access to individual data files.

The site administrator is the warrant of the local site data control.

On each site, as many groups as needed may be created by the system administrator. A group is a unique group identifying name associated to a list of users known by their DN. Note that a group may contain users from different sites (except for the site-specific group as detailed below). A registry service (see next section) will ensure the unicity of group names among different sites.

Two groups will be automatically created on each site upon system installation:

- A site-specific group. All members registered to a site will belong to this group. By default all data registered to a site will be readable by members of the site group. No members of other sites can be registered into the site-specific group.
- An administrator group containing the administrator user. All data registered to a site will be readable and writable to the administrator group. No normal user can be added to this group. Thus, only the administrator has deletion capability over the site data.

Other groups are created and populated without restrictions by the site administrator. The site administrator also grants group-based access to each data file.

2.2. Files sharing across sites

A group is locally administrated by a single site administrator but users from different sites may be registered into a same group. Conversely, site administrators can grant access to their local files for groups owned by external sites. Thus, users belonging to different sites can share data from multiple sites upon joint authorization by the group administrator and the site file administrators. Each site controls the access to its files and the administrator is the warrant of the application of the site access control policy.

2.3. Example

Let us consider as an example the deployment of two sites A and B as illustrated in Figure 1.

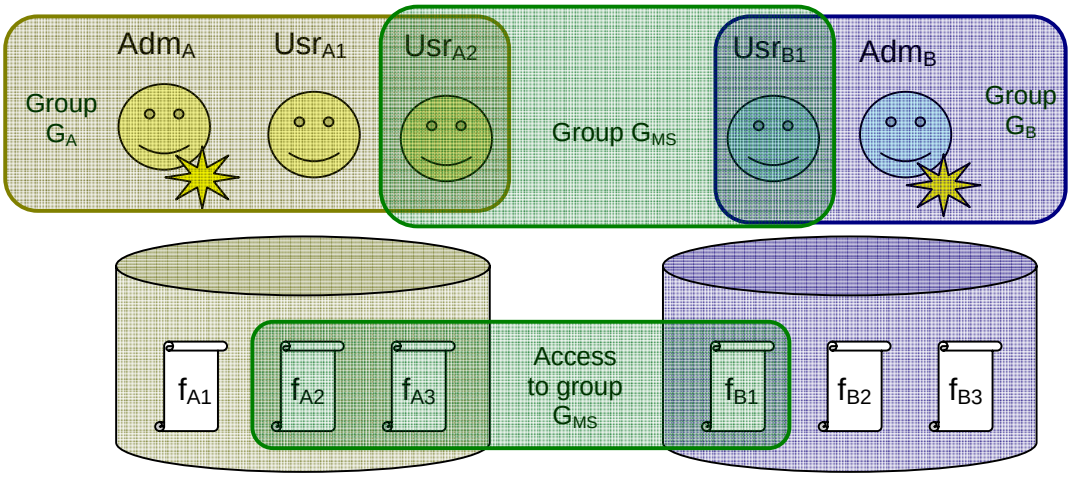

Figure 1. Two-sites deployment example.

Site A is administrated by Adm_A and hosts users Usr_{A1} and Usr_{A2}. Site B is administrated by Adm_B and hosts user Usr_{B1}. Upon deployment, the site specifics groups G_A (containing Adm_A, Usr_{A1} and Usr_{A2}) and G_B (containing Adm_B and Usr_{B1}) are created. Two singleton administrator groups G_{AdmA} and G_{AdmB} are also created.

Suppose know that three files are registered into site A (f_{A1}, f_{A2} and f_{A3}) and into site B (f_{B1}, f_{B2} and f_{B3}). By default, all users of G_A have read access to f_{A1}, f_{A2} and f_{A3} and only G_{AdmA} has deletion capability over these files.

Suppose know that Adm_A creates a group G_{MS} for a working group of users. The DNs of users Usr_{A2} and Usr_{B1} are haded to G_{MS}. Adm_A grants access right to f_{A1} and f_{A2} for group G_{MS} while Adm_B grants access right to f_{B1} for the same group G_{MS}. As a result, both Usr_{A2} and Usr_{B1} are able to access to f_{A1}, f_{A2} and f_{B1}.

3. Operational set up

To implement the NSP described above, the administration services of the NeuroLOG sites will have to cooperate. A NeuroLOG registry will be set up to facilitate this operation. The registry main role will be to register all sites participating to the platform deployed. The registry will be contacted by NeuroLOG services to:

- register a new site;
- discover the participating sites; and
- create new groups.

The registry is a centralized point of failure. Thus the system should depend as little as possible on it. With the proposed solution, it will be needed only upon site and groups creation which are believed to be rare events. In a long term it would make sense to replicate this service to ensure better fault tolerance.

Upon a new site deployment, the registry will be contacted. It will register a unique site name, the site administration service host IP, the site administrator certificate, and the administrator email address. At any time the registry can be queried by one of the sites to discover the other sites registered: thus, the list of sites used e.g. by the Data Federator client can be updated. Furthermore, the registry will allocate a single site prefix to each site, thus ensuring the uniqueness of UIDs generated on each site.

Upon group creation, the site registry will be contacted to register the new group. The registry can be queried at any time to discover existing groups by any participating site.