Integrating CERN e-groups into TWiki access control.

Peter L Jones
IT Department, CERN
Geneva 23 CH 1211
Switzerland
Peter.L.Jones@cern.ch

Alexander Bernegger
IT Department, CERN
Geneva 23 CH 1211
Switzerland
Alexander.Bernegger@cern.ch

Nils Hoymr
IT Department, CERN
Geneva 23 CH 1211
Switzerland
Nils.Hoymr @cern.ch

Abstract

Wikis allow for easy collaborative editing of documents on the web for users located in different buildings, cities or even countries. This is particularly useful for organizations like CERN, hosting globally distributed projects. TWiki (twiki.org) is a project-oriented Wiki implementation, targeting corporate Intranets.

TWiki has been used at CERN since 2003 and has grown popularity and the statistics from December 2009 show over 7000 registered editors and almost 58000 topics (Figure 1). TWiki culture lends to open freeform editing and most pages are world readable and editable by CERN authenticated users, however access control is possible and is used to protect sensitive documents.

The default access control method uses WikiNames or TWikiGroups (topics containing several WikiNames) to define who is allowed or denied access to documents. Since the startup of the LHC more and more groups are requesting finer control for read and write access and it would be more convenient and secure to use the knowledge of user groups from the centrally controlled IT knowledgebase – the E-groups. This note discusses the integration of E-groups for authorisation purposes at CERN.

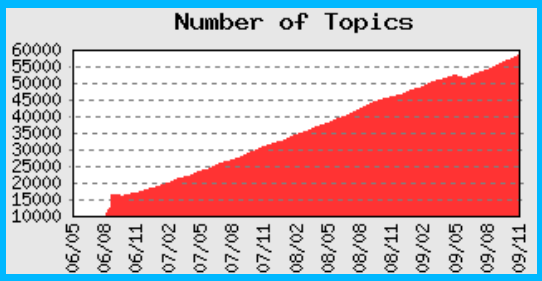

Figure1. Growth of the total number of TWiki files.

1 Introduction

CERN has implemented a Single-Sign-On (SSO) solution based on Microsoft Active Directory Federation Services (ADFS). The SSO allows for centralised authentication to different computing resources and web applications at CERN, thus simplifying user logins and improving application security with a single, secure login infrastructure. Authorisation and access to computing resources is provided via ADFS groups, which are managed via a so-called “E-groups” application that allows users and local administrators to define groups of persons for mailing lists and access to resources. Furthermore, E-groups reflecting organisational groups are generated automatically from data in CERN's HR database. Since 2007 the TWiki instance at CERN has been using SSO authentication.

TWiki topics (known as documents or pages) are contained in areas known as webs and a user can read these TWiki topics from a web browser, which remembers user information for the current session. Many topics are not protected so the user does not have to login.

At CERN if a user wants to edit or read a protected page then the user must authenticate (login) using the CERN SSO account. If not yet registered the user is redirected to the registration page. The user's WikiName and email address are then stored and the TWiki Session retains the user's TWikiName after logon. When a user authenticates the user's credentials are returned by the web daemon that includes AFDS user groups, CERN username and email address.

Access control can be set on whole webs or individual topics. To set access control the user's WikiName or a TWikiGroup is used. A TWikiGroup is made up of WikiNames or other TWikiGroups.

Default access control

Example 1 shows how a web can be configured to allow two TWiki users access to the web.

- Set ALLOWWEBVIEW = TomJones, DickSmith
- Set ALLOWWEBCHANGE = TomJone

Example 1. Access control for users.

The above settings are defined in the web's preferences file and in this example both users Tom and Dick have read access but only Tom has the right to edit topics. If there are many users then a TWikiGroup would facilitate this process. If Tom, Dick are part of a project along with many other users then the following example shows how a TWikiGroup can be used to manage access control

- Set ALLOWWEBVIEW = ProjectOneGroup
- Set ALLOWWEBCHANGE = TomJones

Example 2. Access control using TWikiGroups.

In Example 2 ProjectOneGroup refers to a TWiki topic that contains a list of project members including Tom and Dick. Similar settings can be added inside individual topics and for finer control one can also use DENY settings as well.

So managing access control using the default method is straight forward but the drawback is that the access control lists or TWikiGroups have to be kept up to date and for some experiments the number of users concerned can run into the 1000s.

Implementing access control using E-groups inside TWiki would help collaborations manage their protection policies and therefore increase security.

2 Implementation

An existing TWiki LDAP TWiki extension could help with the implementation, however it is quite a large extension that is no longer maintained. Before looking at this we need to analyze the problem.

Any modification to the TWiki engine to integrate e-groups must coexist with the default access control methods. A grammatical parse of the description of the current system (see introduction) shows the following *classes* and *operations* for a static view.

- *Topic* : getWebDir, getReadProtectedinfo, getWriteProtectedinfo, getRenameProtectedinfo
- *Web* : getReadProtectedinfo, getWriteProtectedinfo, getRenameProtectedinfo
- *User* : getTWikiName, getEmailAddress, getUserGroups, getCERNaccount
- *TWikiGroup* : getTWikiName, getTWikiGroup
- *TWikiSession* : getTWikiName, allowAccess

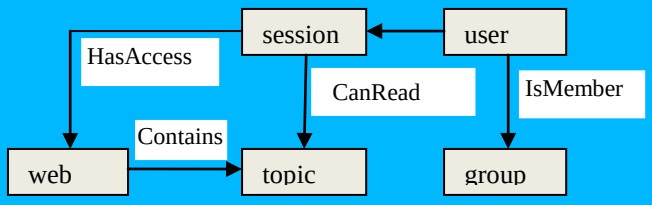

Figure1. Class Diagram of user interaction

The above diagram shows the concept of the problem domain and the inter-relationships of the classes. The following use case scenarios show how TWiki users interact with the system and models the dynamic view of the system.

1. A user wants to read a public topic. No authentication is required and the user has access to the topic as TWikiGuest.
2. A user wants to read a protected topic and is asked to login using the SSO account. The user has access the topic is then displayed. TWiki retains the users TWikiName.
3. A user wants to read a protected topic and is asked to login using the SSO account. The page has some access control and this user is allowed to view the topic. The topic is displayed. TWiki retains the users TWikiName.
4. A user wants to read a protected topic and is asked to login using the SSO account. The page has some access control and this user can not view the topic. An error is displayed. TWiki retains the users TWikiName.
5. A user wants to read a protected topic and is asked to login using the SSO account. The web that contains the topic has some access control and this user is able view pages in this web. The topic is displayed. TWiki retains the users TWikiName.

Taking one of the use cases we can view the interaction between the classes over time. From this candidate operations for the new functionality can be identified.

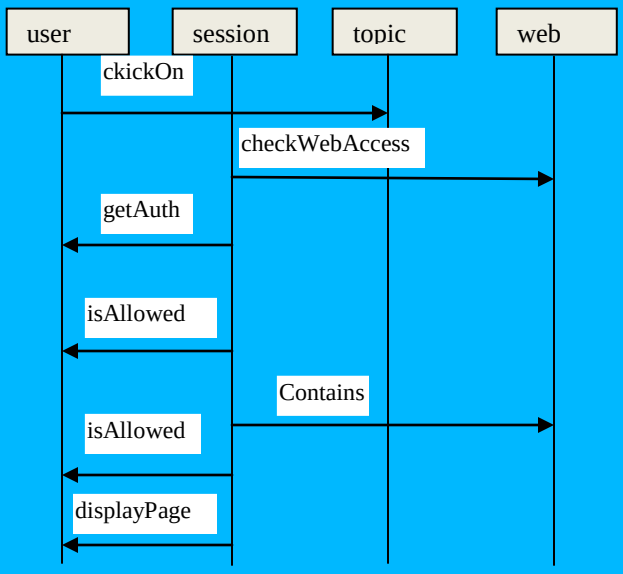

Figure2. Sequence Diagram for Use Case 4.

Re-Use and Inheritance

TWiki is written in Object Oriented Perl and is made up of over 300 separate modules which can be re-used to incorporate any new functionality. From the above we can identify the operators *getAuth* and *isAllowed* that need consideration.

On top of this the Object Oriented nature of TWiki allows us to overwrite existing operations and at the same time inherit the default user mapping module (Figure 3).

```
#
package TWiki::Users::ADFSUserMapping;
use base 'TWiki::Users::TWikiUserMapping';
```

Figure3. Inheritance of original user mapping.

Access control for all DENY and ALLOW options is checked by the TWiki Perl subroutine *checkAccessPermission* from the Access.pm module. Figure 4 shows a snippet of the code that checks the user against ALLOWTOPIC which returns a Boolean.

```
sub checkAccessPermission {
.
.
 # Check ALLOWTOPIC.
 if($users}>isInList($user, $allowText )) {
            return 1;
 }
.
.

}
```

The TWiki configuration is then modified to use the new ADFS user module.

Figure 5 shows an extract of the script that uses the **wget** command that in turn uses the http protocol to create 20 new topics.

Figure 5.Part of the code to check permissions

The sub-routine *isInList* found in the TWiki module Users.pm returns true if the user identification is in a list of user wikinames**,** logins and group ids. This sub-routine is called by *checkAccessPermission* and can be the basis for the implementation of the *isAllowed* operator that was identified earlier.

The apache web server daemon provides useful information about the current user that can be used in the implementation. The environment variable $ENV{HTTP_ADFS_GROUP} contains a list of e-groups that the current user belongs to.

```
sub isInList{
.
.
foreach $ident(split( /\,/,$userlist )) {
.
. $adfs_groups = $ENV{HTTP_ADFS_GROUP};
.
 foreach $adfs_grp(split(/[;]/,$adfs_groups)){
  if ($adfs_groups_list =~ m/^$ident$/i) {
                return 1;
  }
.
.
 }
.
.

}
```

Figure 6.Part of the implementation

Figure 6 shows part of the code that checks the users ADFS information against the entries set in the TWiki access control setting. The variable *$ident* represents each entry. The new version of the sub routines are put in the ADFSUserMapping module.

3 Results

Following the implementation tests were made by setting access control to individual topics and on the level of the web.

Table 1 shows the results of two users that tested out the access control variables ALLOWTOPICVIEW and ALLOWWEBVIEW. User A is a member of the egoup *catia-users* only and User B is only a member of the e-group *service-sdt-user*. Each user attempted to read the topic that was protected by an e-egroup, for example:

- Set ALLOWWEBVIEW = catia-users

Web	**Topic**	**A**	**B**
none	catia-users	√	X
none	service-sdt-user	X	√
none	british-at-cern	X	X
catia-users	none	√	X
service-sdt-user	none	X	√
british-at-cern	none	X	X
british-at-cern	catia-users	√	X

Table1. Results of Access Control on Web and Topic.

4 Discussion

The above tests showed that the access control settings for the variables ALLOWTOPICVIEW and ALLOWWEBVIEW behaved as required. In addition further tests showed that setting a comma delimited list of egroups and TWikiGroups also worked successfully. However further tests show that this implementation does not allow for access control using egroups within TWikiGroups.

5 Conclusions

The above implementation has successfully incorporated ADFS information into the TWiki engine in order to satisfy the requirement for access control using e-groups along with the ALLOWTOPICVIEW and ALLOWWEBVIEW TWiki variables. On top of this a mix of access control settings using e-groups and TWikGroups is also possible. Traditional TWiki access control does allow for fact that a TWikiGroup can contain several other TWikiGroups. However the above solution does not allow for e-groups to be contained within other TWikiGroups and this needs to be addressed in a future cycle.

As TWiki is written in an Object Oriented way it is easy to integrate new code without disturbing default functionality. Basic access control using e-groups as explained above is now possible and is being used by experiments since November 2009. Further work is necessary to use egroups within TWikiGroups.

References and Bibliography

TWiki.org. (2009), http://www.twiki.org/

Jones P. (2009), CERN TWiki statistics, https://twiki.cern.ch/twiki/bin/view/Main/CERNTWikiStatistics

CERN. (European Organization for Nuclear Research) (2009), http://user.web.cern.ch/user/cern.html.

E- groups, (European Organization for Nuclear Research) (2009), https://groups.cern.ch/