

Sakai User Id and EID

June 13, 2006

Please direct questions or comments about this document to:
Glenn R. Golden, ggolden@umich.edu

This document describes User ID and the EID separation in Sakai 2.2

For the last few releases, we've been working to introduce the concept of "internal" and "external" identifiers in Sakai in the User records. We started with a single identifier that had to play both of these roles, and have been tweaking the code to separate these out. For 2.2, we have completed this separation.

Internal and External ID

An "internal" ID is something that is unique within Sakai; something that never changes for the life of the User record; something that we can use to store references to User objects in other parts of Sakai. This is what the `User.getId()` returns. It is always used for inner-Sakai references, and should *never* be used when working with anyone or anything outside of Sakai. For instance, it is NOT the thing that we accept as a user ID when the user logs in; end-users have no clue what their internal Sakai identifier is for their User account. It should not be shown on user interfaces to disambiguate a display of an end user's name. It must not be used when talking to the enterprise systems that back our UserDirectory and Group providers. Id is used for things like storing "created by", user preferences, and group membership, and any other inner-Sakai user references.

An "external" ID is something that is unique within Sakai; something that has meaning to the end-users of Sakai; something that has meaning to the enterprise systems that back our UserDirectory and Group providers; and something that we might expect to see changed within the lifetime of our User objects in Sakai. This is what the `User.getId()` returns. An external ID is also an "enterprise" ID. This is what the end-users know as their login ids. This is a value that we can use to identify our users with the enterprise systems that back our providers (although, there may be other enterprise-assigned identifiers for our end-users). Since an EID can change, it is never used store a reference to a User record in Sakai.

While EID is more appropriate for displaying in a user interface (along with the user name, to disambiguate the user) than ID, we provide a `getDisplayId()` method in the User object specifically for this purpose.

42 **Changes**

43

44 Basically, code that is using `User.getId()` inappropriately needs to be changed to use
45 `User.getEid()` or `User.getDisplayId()`. The `UserDirectoryService` also has two new
46 methods to map ids, in case you don't need the entire `User` object:

47

```
48 String getUserEid(String id) throws UserNotDefinedException
```

```
49 String getUserId(String eid) throws UserNotDefinedException
```

50

51 So, anytime you are displaying something to the end user, odds are very good that you
52 need an EID, not an ID, and should use `getDisplayId()`.

53

54 But don't go overboard and start using EID for ID purposes. If the user interface encodes
55 a user ID, for example, for a selection of users from a list, the ID is appropriate, and
56 required to find the user object.

57

58 The provider code, for both the `UserDirectoryProviders` and `GroupProviders`, needs
59 special attention. These work exclusively with the EID, except for very special cases. So
60 the "ID" sent to a provider from Sakai for authentication or `getUser()` or to find that user's
61 group memberships is an EID, and any "ID" set back from the provider is an EID.

62

63 There are some "decorator" classes that wrap over `User` or some other data structure that
64 has a user ID - and are used for Velocity or JSF interfaces. It might be valuable to add
65 `getEid()` and `getDisplayId()` methods to these decorators to surface the EID and the
66 display ID as well as the ID for UI needs. Related to this are some `User`-like
67 implementations (`Agent`, `Member`, `Participant`) that have a `userId` - and will likely need to
68 surface an EID as well.

69

70 **Not changed**

71

72 A user's "my workspace" site is still using the user ID, not EID. Using the EID here
73 would be a mistake, as this is a user reference that we don't want to have to worry about
74 changing later when the user's EID changes. But the portal URL to the site is based on
75 the site's ID, which is less than ideal. To resolve this problem, the portal will be trained
76 to use user EID as an alternate URL to the user's "MyWorkspace". This is sort of like
77 using a site's alias as well as the site ID in a portal URL to the site (which we can train
78 the portal to do as well).

79

80 A user's "home" content area in `ContentHosting` is still `/user/<user id>`, again because it is
81 a `User` object reference and should not get changed. Most of the user interface already
82 hides this from the end user - but some areas such as the content URLs via `Access` and
83 `Web` and `WebDav` could be further trained (if needed) to use the user's EID as an alias to
84 this content area.

85

86 User references in session and event records are still using user ID. Scripts that query
87 these need to have a way to get at users' actual ids - but since we have a nice table in the
88 database to map ID and EID, this should be easy.

89

90 The "admin" and "postmaster" users are still called that - ID and EID. We could change
91 this, but it's not worth it. To be less fragile, don't use the string "admin" in your code
92 (mostly the test code has this), but use the values defined by the
93 `UserDirectoryService` API:

94

95 `ADMIN_ID`

96 `ADMIN_EID`

97

98 That way, if we do every change the ID or EID of the admin, things will continue to
99 work.

100

101 **Implementation Details**

102

103 We are not keeping a full User record for any user we ever see, as once we thought we
104 might. This is misleading, because for external users (i.e. users defined by the provider),
105 we really don't know anything about them inside Sakai except for their ID (which Sakai
106 allocates) and the EID. So instead, we introduce a mapping table that maps ID and EID.
107 This is the `SAKAI_USER_ID_MAP` table that maps the ID and EID fields. The main
108 `SAKAI_USER` table is unchanged. When we have a record defined inside Sakai, there is
109 a `SAKAI_USER` record and an entry in the `SAKAI_USER_ID_MAP` table. When we
110 have a user who is defined solely by the provider, there's only an entry in the map table,
111 which we make as soon as we first meet this user, by seeing the user's EID.

112

113 **Provider Changes**

114

115 Providers must be changed to work with the User object's EID field, and ignore the ID
116 field.

117

118 We are also enhancing the `UserDirectoryProvider`, to allow it to optionally get involved
119 with the formation of the user `displayId()`. Usually, the `displayId()` is the EID value.
120 There are some special cases where the EID is not the best value to display in user
121 interfaces to disambiguate the user's name display. In these cases, the
122 `UserDirectoryProvider` will implement the new "display advisor" interface, and be called
123 from the `UserDirectoryService` impl. whenever `User.getDisplayId()` is called.

124

125 The provider can use the User object passed in to get the EID, and lookup information
126 based on that, or it can have already stashed away stuff in the User object's properties
127 (either when the provider synthesized the User object, or when a bulk-load process
128 created the internally kept User object), information it will use to form the `displayId`.

129

130 Internal Sakai User objects, and systems without providers or with providers that are not

131 DisplayAdvisors will simply use the EID as the displayId.

132