Kerberos Identity Management (KIM) API

Alexandra Ellwood
MIT Kerberos Consortium
December 11, 2007

What is a Kerberos Identity?

A Kerberos identity is a unique name for identifying an entity to Kerberos servers and services

Also called a “client principal”

Represented as a string: jdoe@EXAMPLE.COM

A user may have more than one Kerberos identity

Multiple Kerberos Identities

Users who have identities with different privileges

jdoe@EXAMPLE.COM and jdoe/admin@EXAMPLE.COM

Users who belong to multiple organizations which do not support cross-realm authentication

jdoe@BANK.COM, jdoe@UNIVERSITY.EDU, etc.

As Kerberos is more widely adopted the number of users with multiple identities will increase

Goals of the KIM API

- Help applications select a Kerberos identity and acquire initial credentials if needed
- Cross-platform, object oriented and extensible
- A consistent and friendly user experience
- Support for all Kerberos standards and site configurations

User Experience Today

Users manually switch between identities

Command line user experience:

- Choose identities by setting the KRB5CCNAME environment variable
- Identity inherited with the process environment
- Can use different identities in different terminal windows

User Experience Today (cont.)

Graphical user experience:

- Choose the default identity with a graphical application
- Switching between identities affects all applications

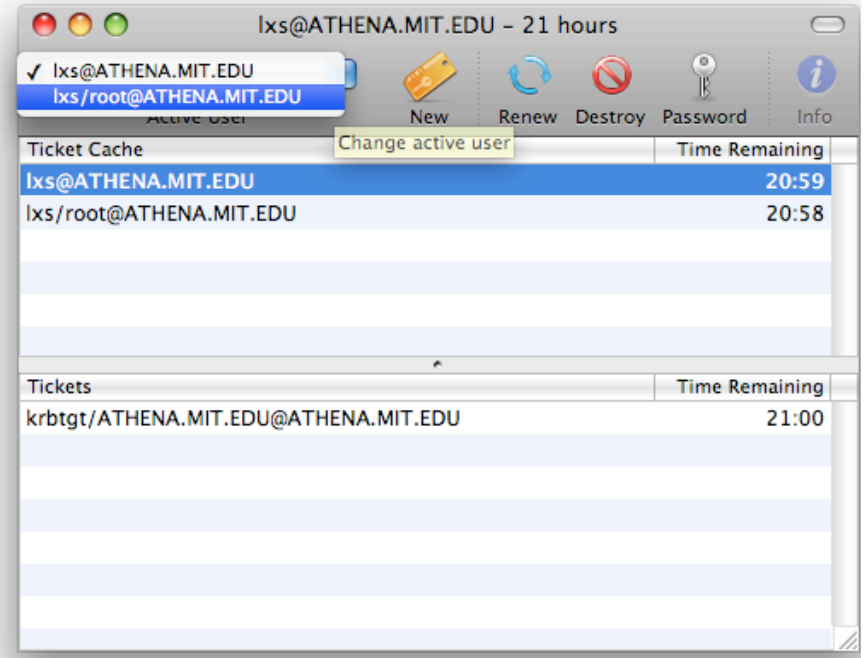

- If applications use different identities only one may work at a time -- others may prompt for new credentials even if valid credentials are available but not default

Improving User Experience

Trying each identity

- May be unacceptable on slow networks or mobile devices
- Exposes the names of the services a user uses to each organization with which they have an identity
- May use the wrong identity if multiple identities can connect to the service

  (eg: using jdoe/admin@EXAMPLE.COM instead of jdoe@EXAMPLE.COM)

Improving User Experience (cont.)

Guessing the identity not supported by standards

- Cross-realm authentication -- the identity and the service may not be in the same realm

- Server referrals -- need to know the identity to determine the actual name of the service

- Choosing between multiple identities in the same realm requires knowing the organization's naming policy

(jdoe/admin@EXAMPLE.COM versus jdoe/root@EXAMPLE.COM)

KIM Identity Selection

- Allows an application to prompt the user once and then look up which identity worked
- Uses mapping from “selection hints” to identity:

  (application name, server hostname, account name, etc)
- Support application protocols which can determine the identity and avoid prompting the user
- Support for future Kerberos extensions that make guessing the identity more feasible

KIM User Experience

Prompt the user to select an identity

- Presents a list of favorite identities

- Offers the option of getting credentials for a new identity

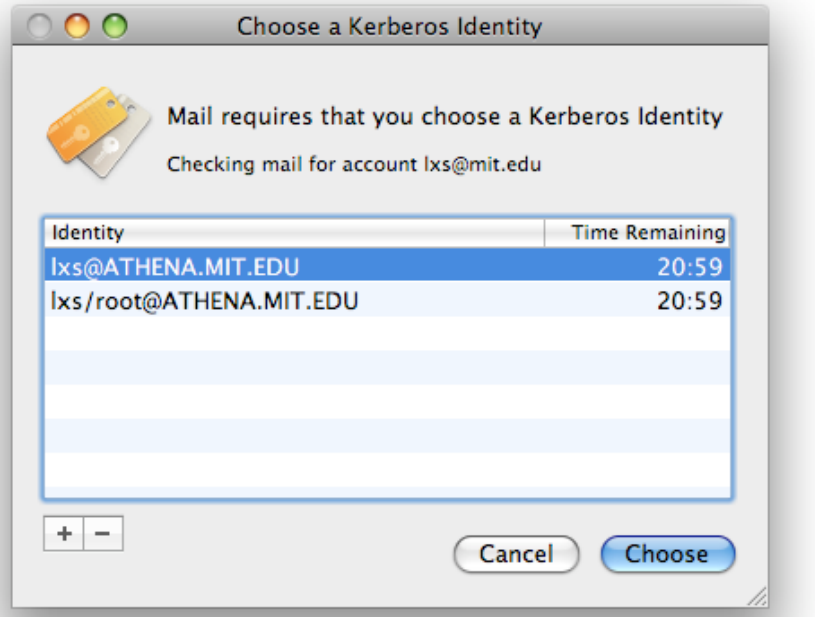

- Application-provided localized strings explain why the application needs credentials

Other KIM Features

APIs for creating and destroying initial credentials (TGTs) and credentials caches

(alternative to krb5_init_creds and krb5_cc_* functions)

Per-user preferences for ticket options

Plug-in APIs for handling credential acquisition and destruction

Documentation

KIM API Draft:

http://mit.edu/macdev/kim.html

Send feedback to krbdev@mit.edu