

9-16-11 Progress Report

What we did

Based on its initial charge, the Registries subteam developed a functional model of Identity Registries, a Glossary, and collected a list of functional requirements from subteam participants. The requirements list is fairly comprehensive in that it covers almost areas of interest in a Registries system, but the list has not been edited for consistency, nor prioritized. More requirements could be added.

Three projects assessed their capabilities against the requirements: OpenRegistry, the PSU Central Person Registry, and Kuali Identity Management. There are pluses and minuses to each projects, as detailed below. A fourth project, ForgeRock, may also be assessed. At this point no one project stands out as the obvious one to choose based on either current functionality or likely future wonderfulness.

OpenRegistry

OpenRegistry is a system designed to be very close to the functionality the subteam has defined for a Registries system. It is intended to be a system that could be deployed at many sites.

OpenRegistry meets most of the collected Registries requirements.

OpenRegistry is still under development and has not yet been deployed by any institution (yes?), so there is no information about how well it performs in practice, or how well it meets needs for flexibility and extensibility. The project historically has suffered from some lack of continuity and lack of participation. It is currently primarily driven by one institution. Some institutions have used its code (i.e., forked its code) as a base for their own registry development.

PSU Central Person Registry

The PSU-CPR is also a system designed to be very close to the functionality the subteam has defined for a Registries system. PSU-CPR has been developed by PSU to meet its institutional needs, which are similar to those of other large institutions.

PSU-CPR meets most of the collected Registries requirements.

PSU-CPR has not been specifically designed to be used by other sites; e.g. no external licensing or distribution packaging. It is still under development, so there is no information about how well it performs in practice, or how well it meets needs for flexibility and extensibility. The PSU project is funded and staffed to be completed and in production by the end of 2011.

Kuali Identity Management

KIM, a component of Kuali Rice provides a large set of IAM services; for this purpose the subteam is looking at KIM's identity registry functions. KIM is designed to meet the IAM needs of Kuali applications.

KIM meets most of the collected Registries requirements, though somewhat fewer than OpenRegistry or PSU-CPR.

KIM has been in use by many sites in production for a few years. Version 2.0 is under development and is slated for release in Q1 2012. Extending the current system to meet institutional needs as defined by the subteam would require some redesign. Kuali Rice has an established means for organizing interested parties to support large development projects.

ForgeRock

ForgeRock is a commercial operation developing and supporting an open-source set of IAM packages many of which are based on the open-source Sun packages that made up the Sun IAM product suite. The products have been reasonably successful and the company has grown in the last couple of years. There may be opportunities for engagement from our community.

The ForgeRock suite does not appear to include a component that matches the functionality defined by the Registries subteam.

Bottom line

The team has general agreement on these things:

1. An Identity Registry system meeting the identified institutional requirements and able to be deployed at many sites could be developed in a reasonable time (12-18 months?) and reasonable effort (a few FTE?). The three evaluated projects are existence proofs of systems very like what is needed.
2. There are some differences among the participants. Kuali sees identity merge for employees and students as the most compelling feature. The more central-IdM-centric folks see more advanced features like assurance-tracking and schema agility as most important.
3. More work is needed to develop a more concrete path forward. Questions that could be investigated are:
 - a. Can the OpenRegistry design form the basis of the desired system? Evidence is mixed based on institutions that have looked at it in the past.
 - b. Can a more complete and well-articulated set of requirements be put together to entice investors? How long will that take?
 - c. Should OSIdM4HE attempt to choose among projects as a starting point, or leave it to the market to decide?
 - d. How do we really assess the market for this, and relative importance of features, beyond the core group?