Dealing with Multiple Attribute Stores and the Shib IdP

Dealing with Multiple Attribute Stores and the Shib IdP

DATE and TIME: Thurs. May 26, 2011, 10AM

CONVENER: Mike Wiseman, University of Toronto

SCRIBE: Pete St. Onge, University of Toronto

# of ATTENDEES: 20

MAIN ISSUES DISCUSSED

The identified issue is how to handle the use of multiple (more than two or three) data source connections to a shib IdP. In an institution, there may be many data stores managed centrally or in departments. If a department wants their attribute store to be used in shib authorization, should the store be connected to the IdP or should it be associated with the application side. A technical issue noted with connecting multiple data sources to an IdP is that those sources get accessed for each IdP transaction, regardless of the attribute filter settings.

Suggested ideas at the session can be frouped into two categories - IdP-centric and datastore-centric.

IdP-centric:

1. Examine the use of 'attribute aggregation' in shib. This involves placing IdPs in front of data stores. A user acess transaction then involves two or more IdPs.

2. Look at 'virtual directories'. Perhaps these products provide mechanisms to handle the continuous requests from the IdP.

Datastore-centric:

1. Use a product such as Grouper to aggregate adtasources to one physical store.

ACTIVITIES GOING FORWARD / NEXT STEPS

1. Document the use of attribute aggregation.

2. Get input on the multi-datastore handling by the IdP from IdP developers.