Improving FERPA-protected Data Access

As we move towards migrating our Shibboleth IdPs to Amazon Web Services, we have an even greater need to remove its dependency on multiple LDAPs and consolidate all necessary user data to the Active Directory. In addition, as we implement Grouper as the central authorization management tool, we need to ensure that PII is available for authorized applications and units to consume, in order to provide the same user experience regardless of data suppression. Because our AD is a critical component of so many services across campus, it made sense to provide greater functionality and greater security for the good of our campus.

The University of Illinois at Urbana-Champaign has historically maintained two LDAP instances that served slightly different flavors of directory information. One LDAP instance (what we call the Campus LDAP) contains all user attributes with access restrictions on sensitive data classes, while our Active Directory LDAP contains somewhat of a "redacted" identity for those users with FERPA-protected attributes. These redacted identities in AD have no personally identifiable information (PII); this was accomplished by suppressing several personal name attributes (givenName, sn, etc.) in AD by replacing them with null values or a NetID (our name for the logon ID). This has proven to be a poor experience for both the user as well as for units and applications that provide IT services to those users.

In order to eliminate duplication and also provide the best of both LDAPs, our Identity and Access Management team decided to keep the Active Directory and retire the Campus LDAP. This meant that we would need to converge all of the existing attributes and ensure that sensitive data on FERPA-protected users was readily available in the AD, without compromising access controls on those attributes. The solution we arrived at was to create a set of "shadow" attributes that contained all of the PII, while placing them in a custom property set, protected with privileged access. This solution perpetuates the redaction of PII from default attributes that are accessible to the general authenticated public, while providing service accounts and applications the necessary privileged access to the real data. It also requires zero touch by current services-- they can continue to consume the same data as they are used to now without any changes, or they can opt to apply for the privilege to access the protected attributes.