Infrastructure - Task #1598

Story # 312 (Closed): Identifier generation and reservation

Design replicated store for reserved identifiers

2011-05-31 15:15 - Dave Vieglais

Status:	Closed		**Start date:**	2011-05-31
Priority:	Normal		**Due date:**	
Assignee:	Ben Leinfelder		**% Done:**	100%
Category:			**Estimated time:**	0.00 hour
Target version:	Sprint-2012.03-Block.1.2			
Milestone:	CCI-0.6.2		**Story Points:**	
Product Version:	*			

Description
We need a data store replicated across the CNs for keeping track of identifiers that have been reserved.
Expect that the postgres database would be the simplest mechanism for this - using either metacat as the replication manager or some other lock mechanism.

History

#1 - 2011-05-31 15:20 - Dave Vieglais

Note that we will likely use the http://n2t.net/ezid EZID service for identifier production. Some additional considerations listed in #1599

#2 - 2011-07-13 20:43 - Ben Leinfelder

- Status changed from New to Closed

- Milestone set to CCI-0.6.2

We now have an LDAP implementation for reserving Identifiers for a particular user for a short amount of time (24 hour expiration).
The LDAP objectClass hass been added to the cn_buildout and the ReserveIdentifierService class has been coded in the d1_identity_manager project (which shares many similar LDAP functions)