

Infrastructure - Task #2069

Story # 2067 (Closed): Refactor the index task generator and processor to improve scalability and performance

Select a persistence technology for the index task queue

2011-11-29 14:04 - Dave Vieglais

Status:	Closed	Start date:	2011-11-29
Priority:	Normal	Due date:	
Assignee:	Skye Roseboom	% Done:	100%
Category:	d1_indexer	Estimated time:	0.00 hour
Target version:	Sprint-2012.05-Block.1.3	Story Points:	
Milestone:	CCI-1.0.0		
Product Version:	*		

Description

Spring Data (<http://www.springsource.org/spring-data>) seems like a good choice, perhaps with the JPA or JDBC connection. Postgres is installed on the CNs for use by Metacat, would make sense to use this for the persistence mechanism, though need to check on resource use by Metacat to ensure there are sufficient connections available for the indexer to operate.

History

#1 - 2011-11-30 00:39 - Skye Roseboom

- Status changed from New to Closed

Implemented JPA using Hibernate with Spring configuration. Spring Data is providing a JPA Repository so no need to write a DAO.