

Infrastructure - Task #2015

Story # 1744 (Closed): Design and Implement log aggregation capability

Design the log aggregation functionality

2011-11-10 16:22 - Dave Vieglais

<b>Status:</b>	Closed	<b>Start date:</b>	2011-11-10
<b>Priority:</b>	Normal	<b>Due date:</b>	2011-11-14
<b>Assignee:</b>	Dave Vieglais	<b>% Done:</b>	100%
<b>Category:</b>	d1_cn_service	<b>Estimated time:</b>	2.00 hours
<b>Target version:</b>	Sprint-2011.49-Block.6	<b>Story Points:</b>	
<b>Milestone:</b>	CCI-1.0.0		
<b>Product Version:</b>	*		
<b>Description</b>			
Identify the technologies to be used and components that need to be implemented to support the log aggregation service.			
Draft design document capturing the key concepts and layout.			
Note due date.			

History

#1 - 2011-12-02 02:31 - Robert Waltz

- Status changed from New to In Progress
- Assignee changed from Robert Waltz to Dave Vieglais

Please review

<http://mule1.dataone.org/ArchitectureDocs-current/design/LogAggregator.html>

#2 - 2011-12-06 14:42 - Dave Vieglais

There is some concern over the number of messages that need to be transmitted to ensure log replication. The implementation should enable the option of batching log record processing, so that some analysis of optimal message count vs message size can be made.

It should also be noted that adding new records to SOLR is significantly more efficient if done in batches rather than individually.

#3 - 2011-12-13 00:30 - Dave Vieglais

- Status changed from In Progress to Closed

Design appears suitable. Closing this task. Possible iterations of design and implementation may be required with further testing.