

Infrastructure - Story #8084

Story # 8061 (New): develop queue-based processing system for the CN

determine the backup strategy for rabbitMQ

2017-04-26 17:44 - Rob Nahf

Status:	New	Start date:	2017-04-26
Priority:	Normal	Due date:	
Assignee:	Rob Nahf	% Done:	0%
Category:	d1_indexer	Estimated time:	0.00 hour
Target version:			
Story Points:			
Description			
Do we need to back up not only the rabbitMQ configuration, but also messages? Consider that we might have long-lived dead-letter or task failed queues...			

History

#1 - 2017-04-26 20:11 - Rob Nahf

possibility? <https://lists.rabbitmq.com/pipermail/rabbitmq-discuss/2013-August/029124.html>

#2 - 2017-05-09 13:58 - Dave Vieglais

The rabbitmq queues should be replicated across machines, so failure of one or two machines could be restored from the remaining. This is the case with much of the CN infrastructure.

Second, the information contained in the queues could be restored by reindexing content that has been added or modified since time before the failure.

Hence it seems that it should be sufficient to ensure that the rabbitmq configuration is preserved to enable the instances to be reinstalled as necessary.

#3 - 2018-01-17 18:46 - Dave Vieglais

- Sprint set to Infrastructure backlog