

Infrastructure - Story #8081

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

develop federated broker configuration for indexing

2017-04-24 22:52 - Rob Nahf

Status:	In Progress	Start date:	
Priority:	Normal	Due date:	
Assignee:	Rob Nahf	% Done:	30%
Category:	d1_indexer	Estimated time:	0.00 hour
Target version:			
Story Points:			
Description			

History

#1 - 2017-04-25 17:04 - Rob Nahf

from RMQ's networking page: (<https://www.rabbitmq.com/networking.html>)

looks like we need to open ports 5672 and 4369; and possibly 15672 and 25672

Port Access

SELinux, and similar mechanisms may prevent RabbitMQ from binding to a port. When that happens, RabbitMQ will fail to start. Firewalls can prevent nodes and CLI tools from communicating with each other. Make sure the following ports can be opened:

- 4369: epmd, a peer discovery service used by RabbitMQ nodes and CLI tools
- 5672, 5671: used by AMQP 0-9-1 and 1.0 clients without and with TLS
- 25672: used by Erlang distribution for inter-node and CLI tools communication and is allocated from a dynamic range (limited to a single port by default, computed as AMQP port + 20000). See networking guide for details.
- 15672: HTTP API clients and rabbitmqadmin (only if the management plugin is enabled)
- 61613, 61614: STOMP clients without and with TLS (only if the STOMP plugin is enabled)
- 1883, 8883: MQTT clients without and with TLS, if the MQTT plugin is enabled
- 15674: STOMP-over-WebSockets clients (only if the Web STOMP plugin is enabled)
- 15675: MQTT-over-WebSockets clients (only if the Web MQTT plugin is enabled)

It is possible to configure RabbitMQ to use different ports and specific network interfaces.

#2 - 2017-04-26 04:57 - Rob Nahf

- Subject changed from develop federated brokers for indexing to develop federated broker configuration for indexing

#3 - 2017-04-26 05:41 - Rob Nahf

- Status changed from New to In Progress

- % Done changed from 0 to 30

I got simple federated queueing working by setting up firewall rules between servers on ports 5672 and 4369 (administrative channel) and configuring 'upstreams' and policies.: A consumer on one machine can consume messages published on the other machine.

Still need to set up a different account, and maybe add some security, although maybe a named user for the process ("indexer" user) would be sufficient, since the ports are limited. Also figure out how to put the configuration in debian configuration. I stumbled across a rabbitmq.config that defines federated-queues.

<https://github.com/jamescarr/rabbitmq-federation-example/blob/master/node1/rabbitmq.config>

#4 - 2018-01-17 18:46 - Dave Viegla

- Sprint set to Infrastructure backlog