


Brett Jones, Scott Luxenberg, David McGrath, Paul Trampert, and Jonathon Weldon  
 CS 4284 Systems and Networking Capstone Spring 2011 Faculty Advisor: Dr. Ali Butt

## Background

### What is RabbitMQ?

- Application messenger
- Based on publish-subscribe model
  - Content creators (publishers) send content to RabbitMQ messengers
  - Messengers distribute content to content readers (subscribers)
- Advantages:
  - No polling
  - Publishers abstracted from subscribers
  - Publishers not overloaded with pushing data



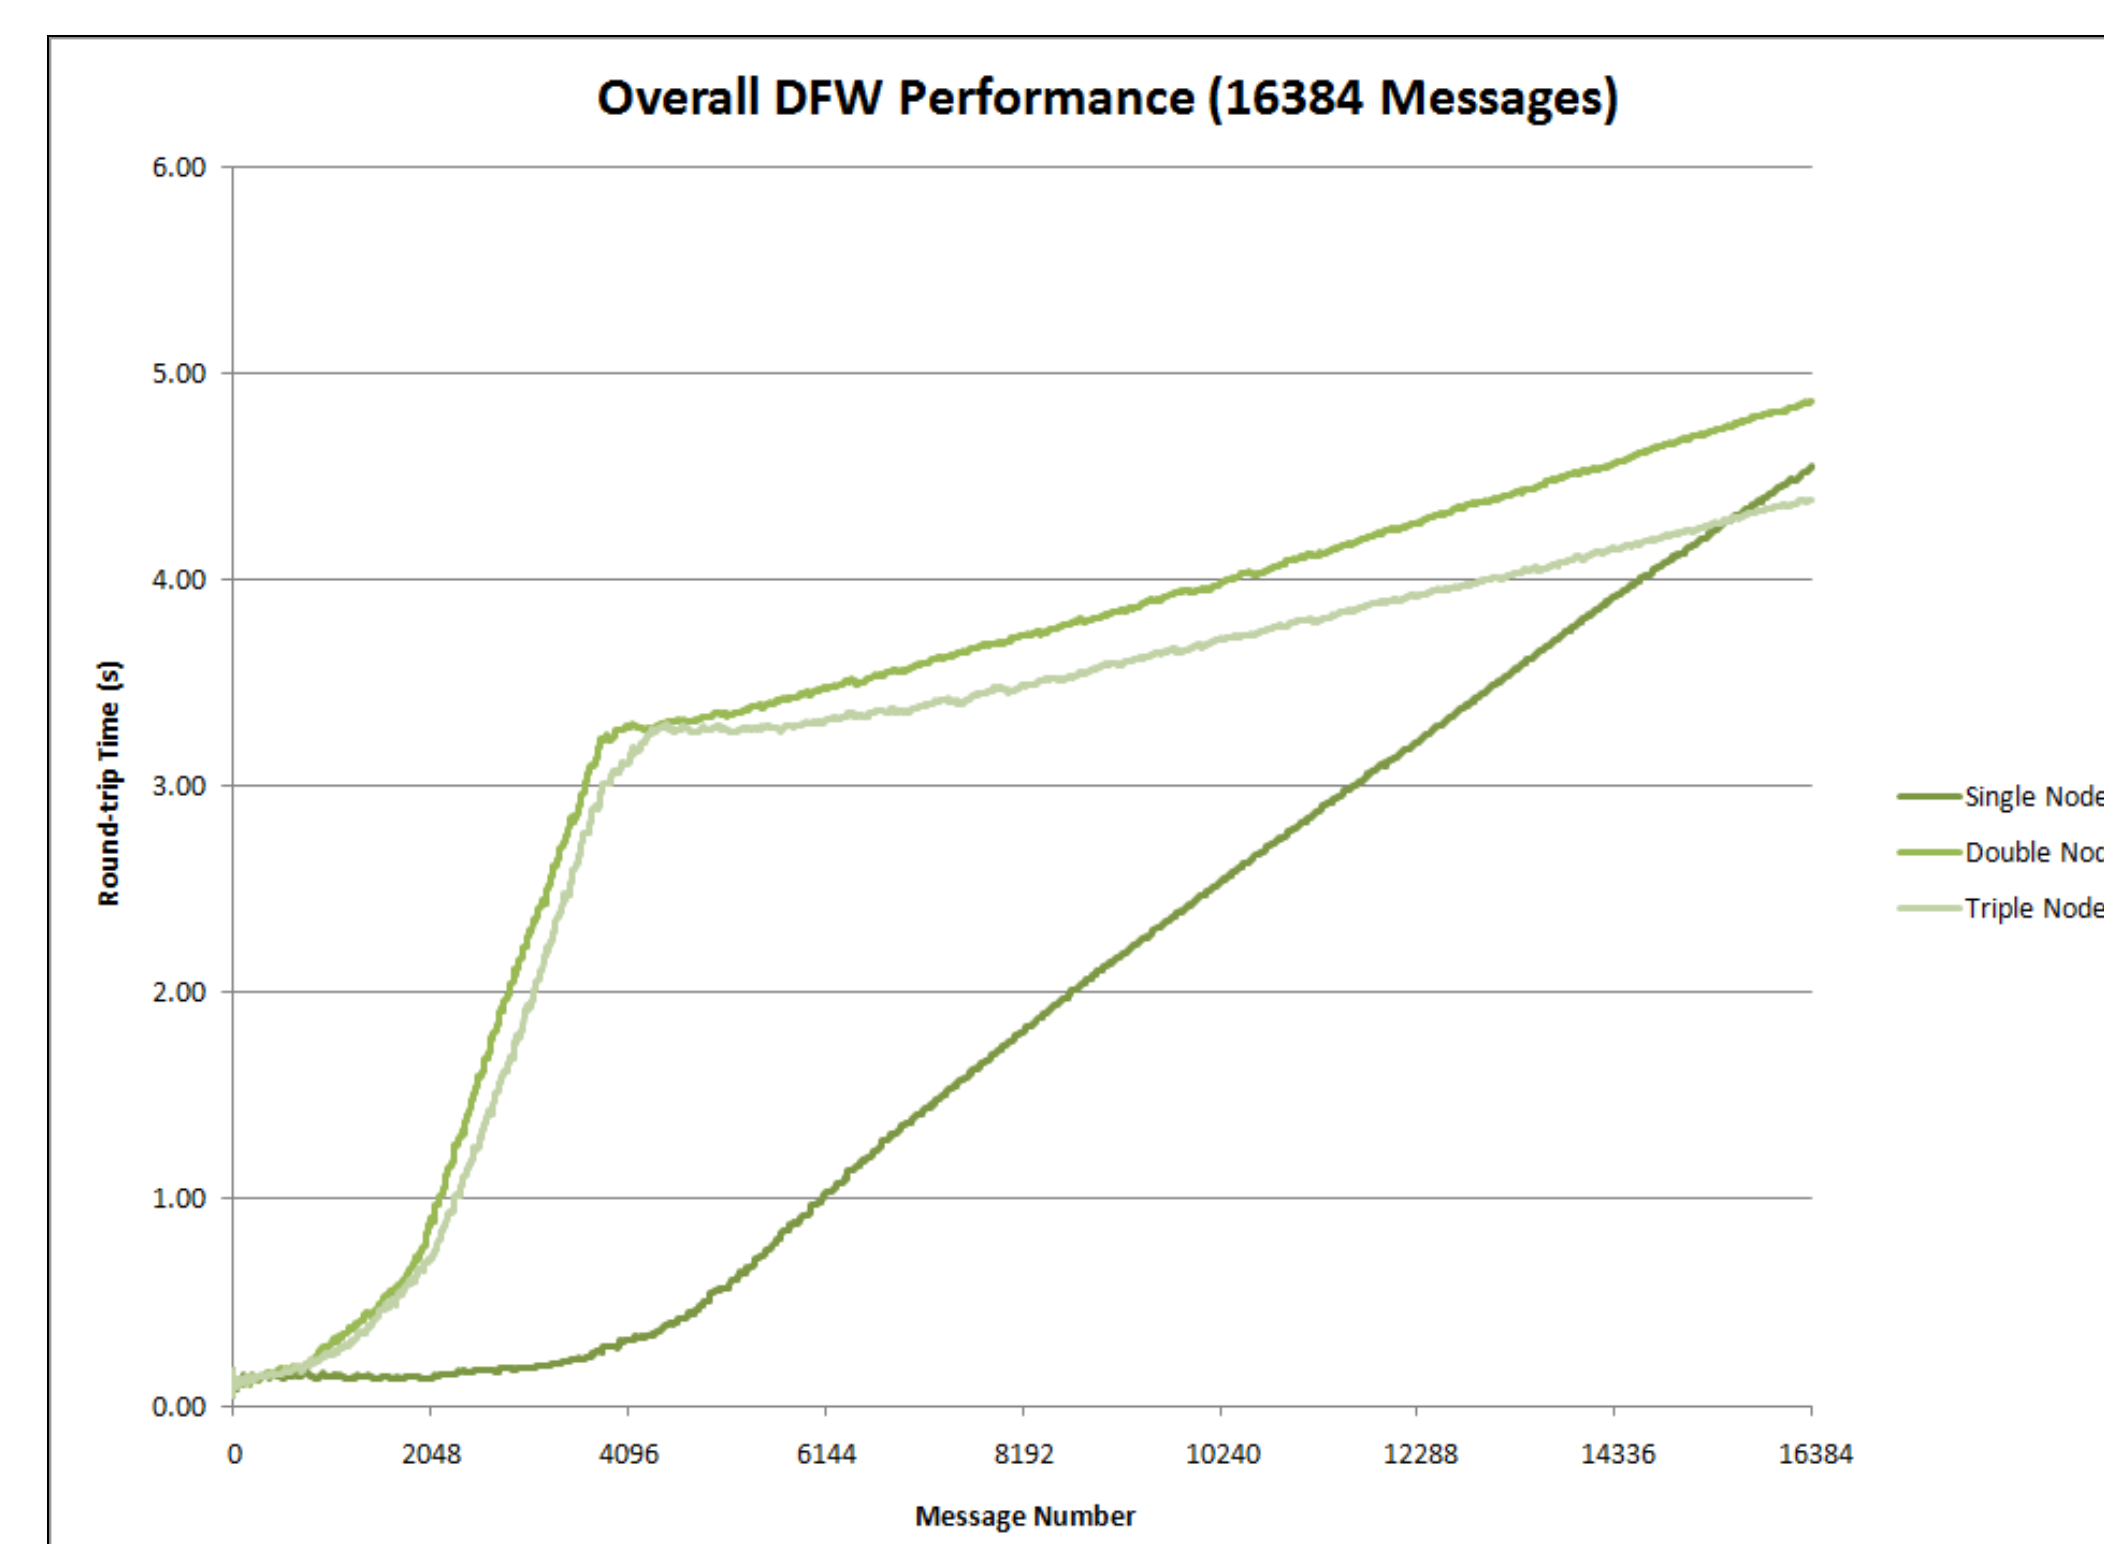

### Motivation

- Rackspace looking for scalable messaging architecture
  - Purposed for distributed system messaging backbone within data centers
- Determining the scalability of RabbitMQ can open up more options for its usage in large-scale applications

## Results

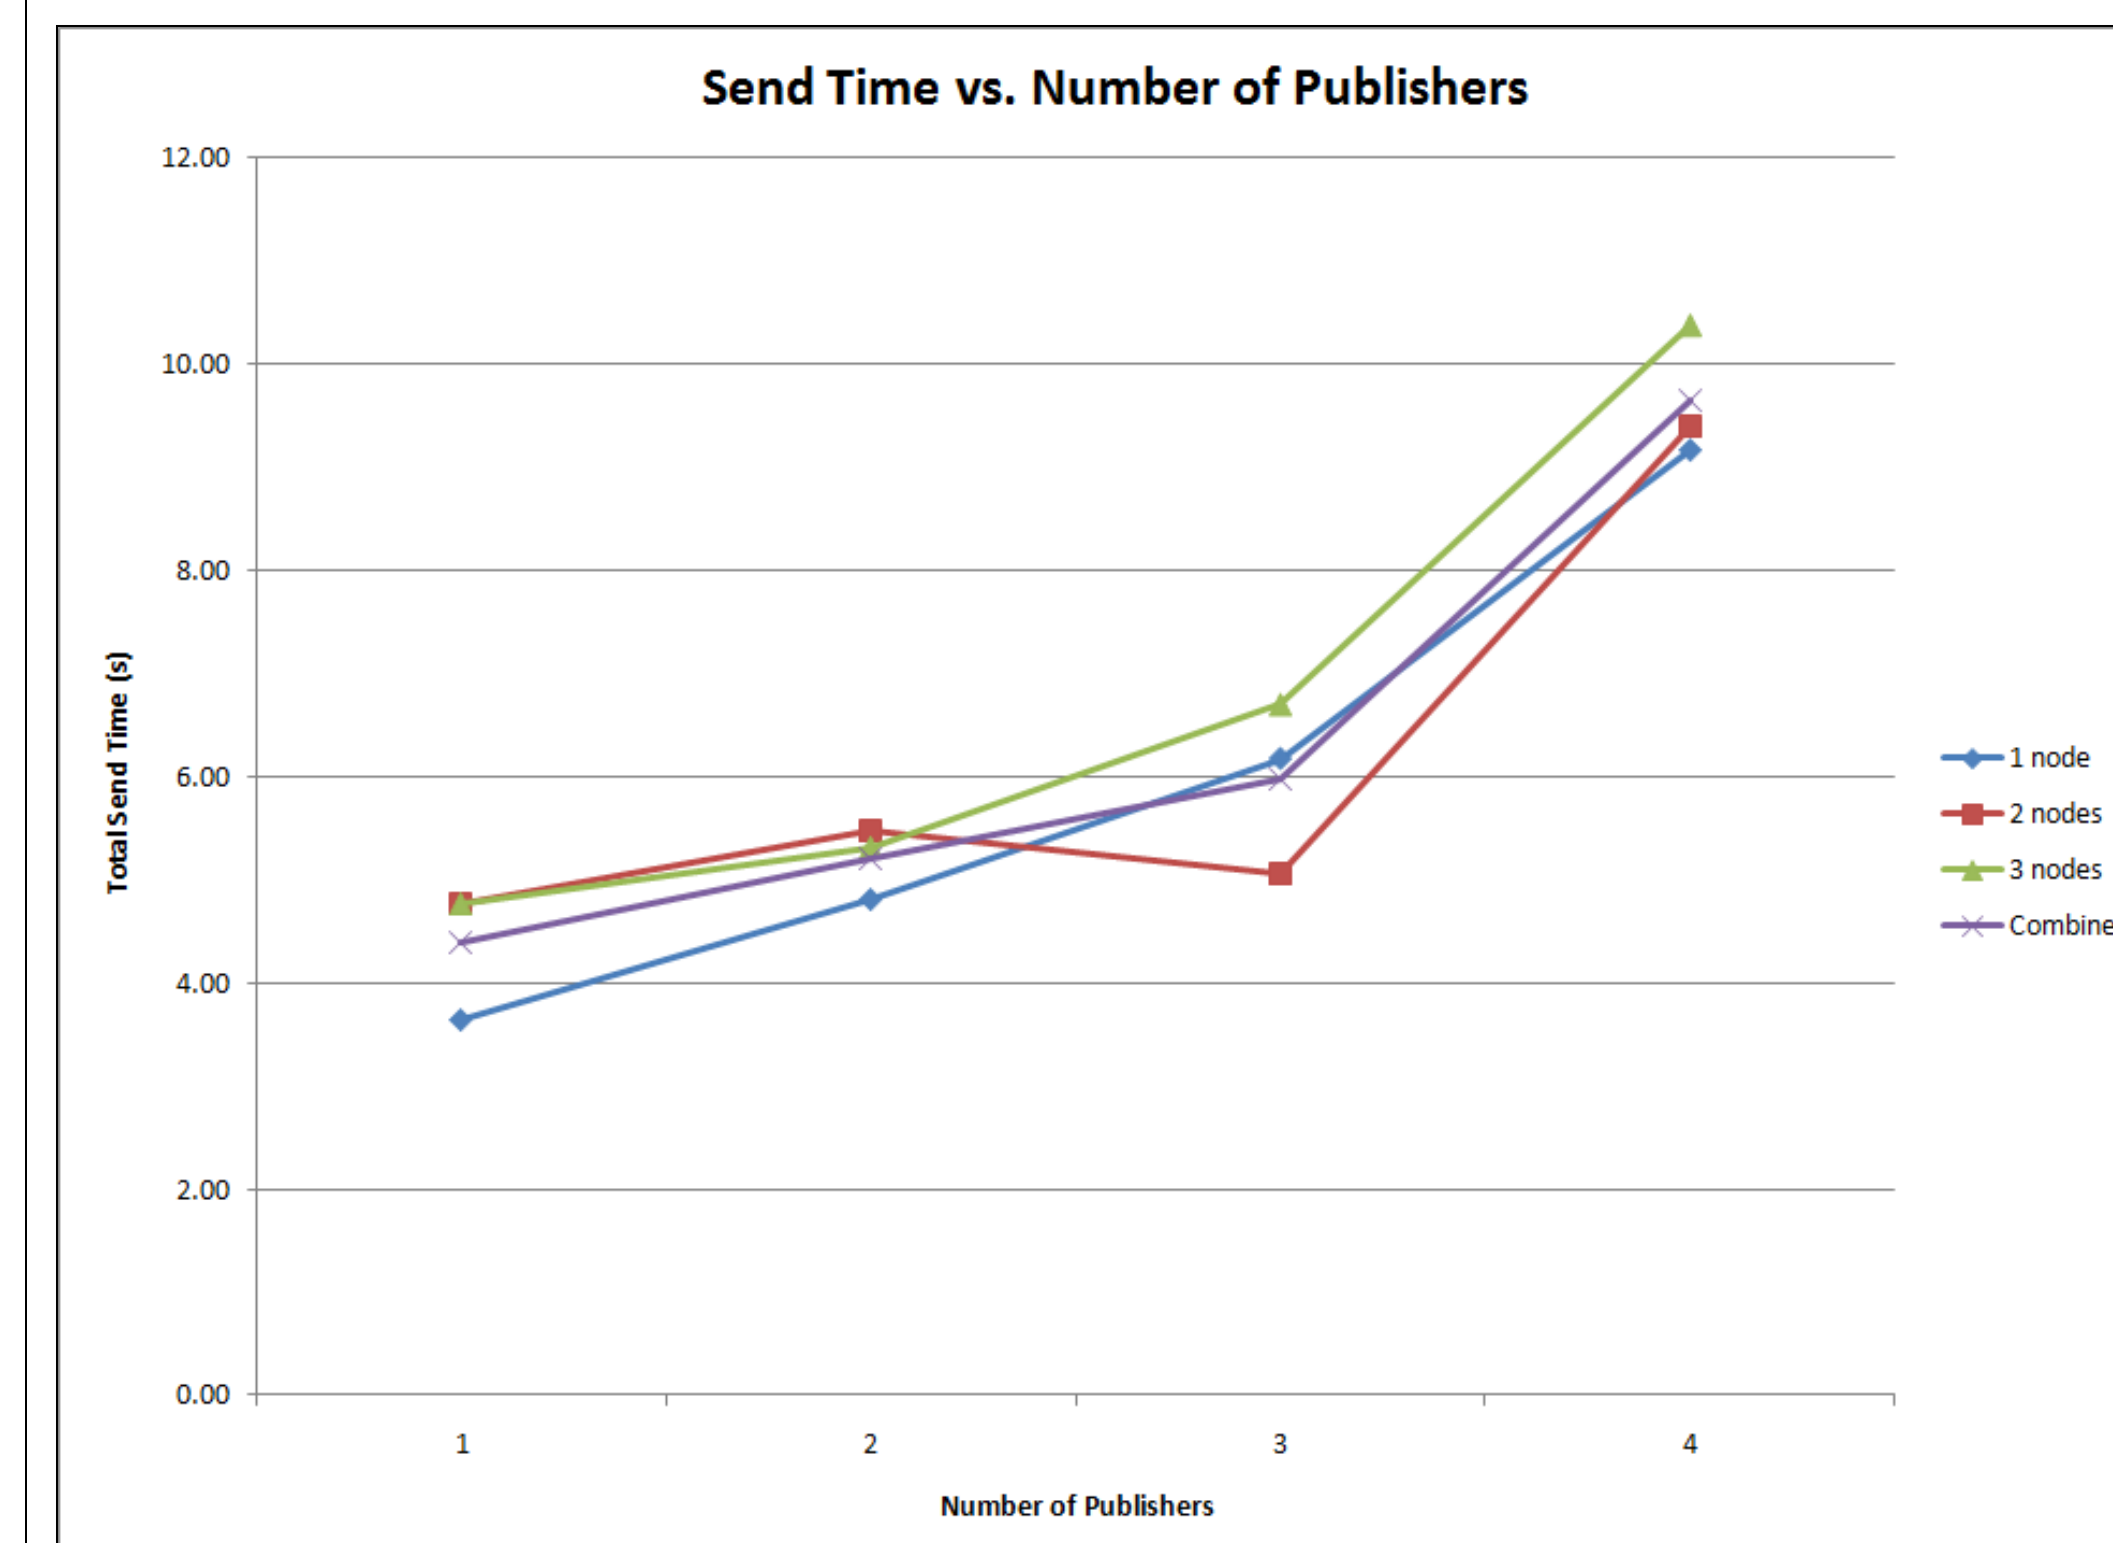
### Single Publisher, Single Subscriber

- Process
  - Tested with 1KB messages
  - Graphed round-trip times for each message
  - Used 4K and 16K message sets with clusters of 1-3 nodes in Dallas/Fort Worth data center
- Results
  - Multi-node clusters initially perform worse than single-node clusters, but showed smaller slopes after receiving all messages



### Multiple Publishers, Single Subscriber

- Process
  - Tested with 1KB messages
  - Graphed times to send all messages from 1-4 publishers
  - Used 16K message sets with clusters of 1-3 nodes in Dallas/Fort Worth data center
- Results
  - Uniform scaling from 1 to 4 publishers
  - No noticeable impact for number of nodes in the cluster



### Single Publisher, Multiple Subscribers

- Plan
  - Tested with 1KB messages
  - Graphed times to receive all messages in 1-4 subscribers
  - Used 8-2K message sets (in powers of 2) with clusters of 1-3 nodes in Dallas/Fort Worth data center
- Results
  - Ended up scrapping this test due to complications with uneven message distribution

### Multiple Publishers, Multiple Subscribers

- Process
  - Tested with 1KB messages
  - Graphed round-trip times for each message
  - Used 16K message sets with clusters of 1-3 nodes in Dallas/Fort Worth data center
  - Each publisher sent messages to a different node
- Results
  - Increasing publishers and subscribers puts highest strain on head node
  - Some messages may be received out of order
    - Exchange vs. queue

