

MongoDB Evaluation

Luke Miller and Alex Wu

Goals

- Deploy MongoDB on Amazon EC2 Instances
- Create a Web Server with Node.js
- Evaluate Performance of Vertical and Horizontal Scaling

Testing Environment

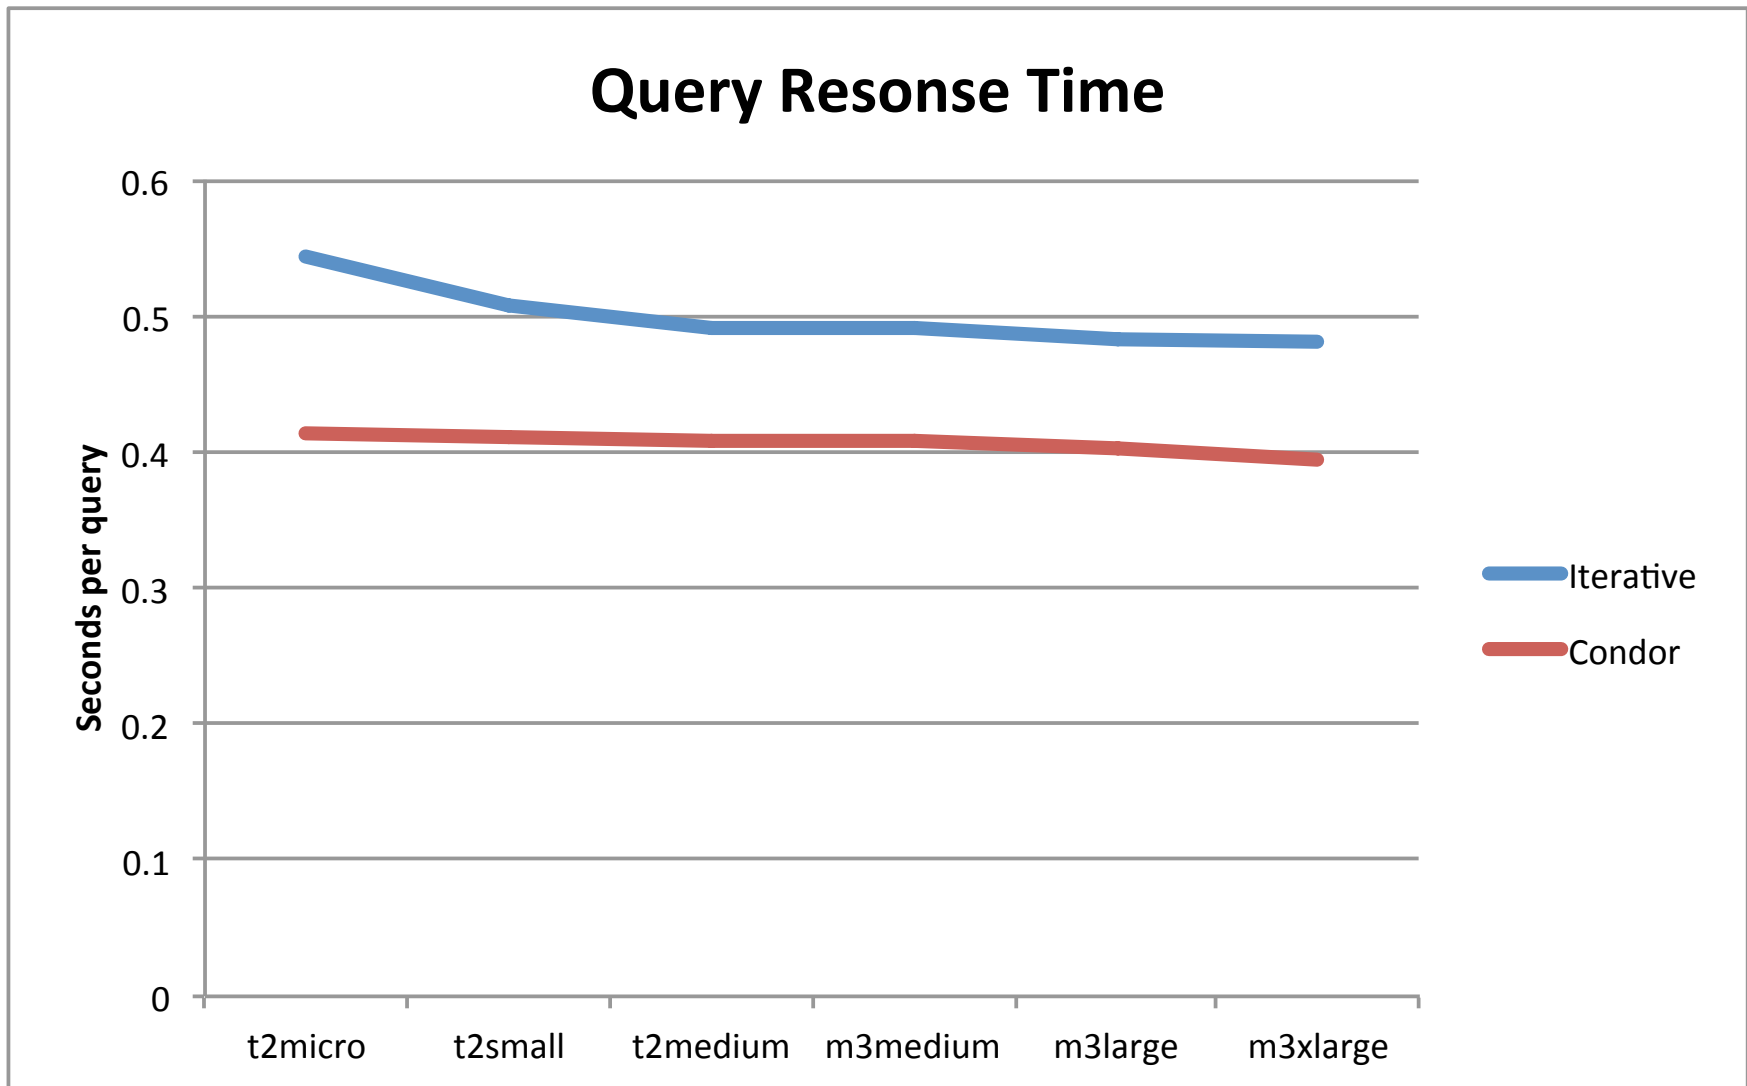
- Data: Over 500,000 tweets from the Boston Marathon Bombing dataset provided by Prof. Wang from the Apollo Project
- Submit Jobs using Notre Dame's Condor Pool to test performance with many "users"

Vertical Scaling

- 1000 Semi-Random Queries
- Iteratively vs Condor Pool
- Comparing EC2 Instance Types

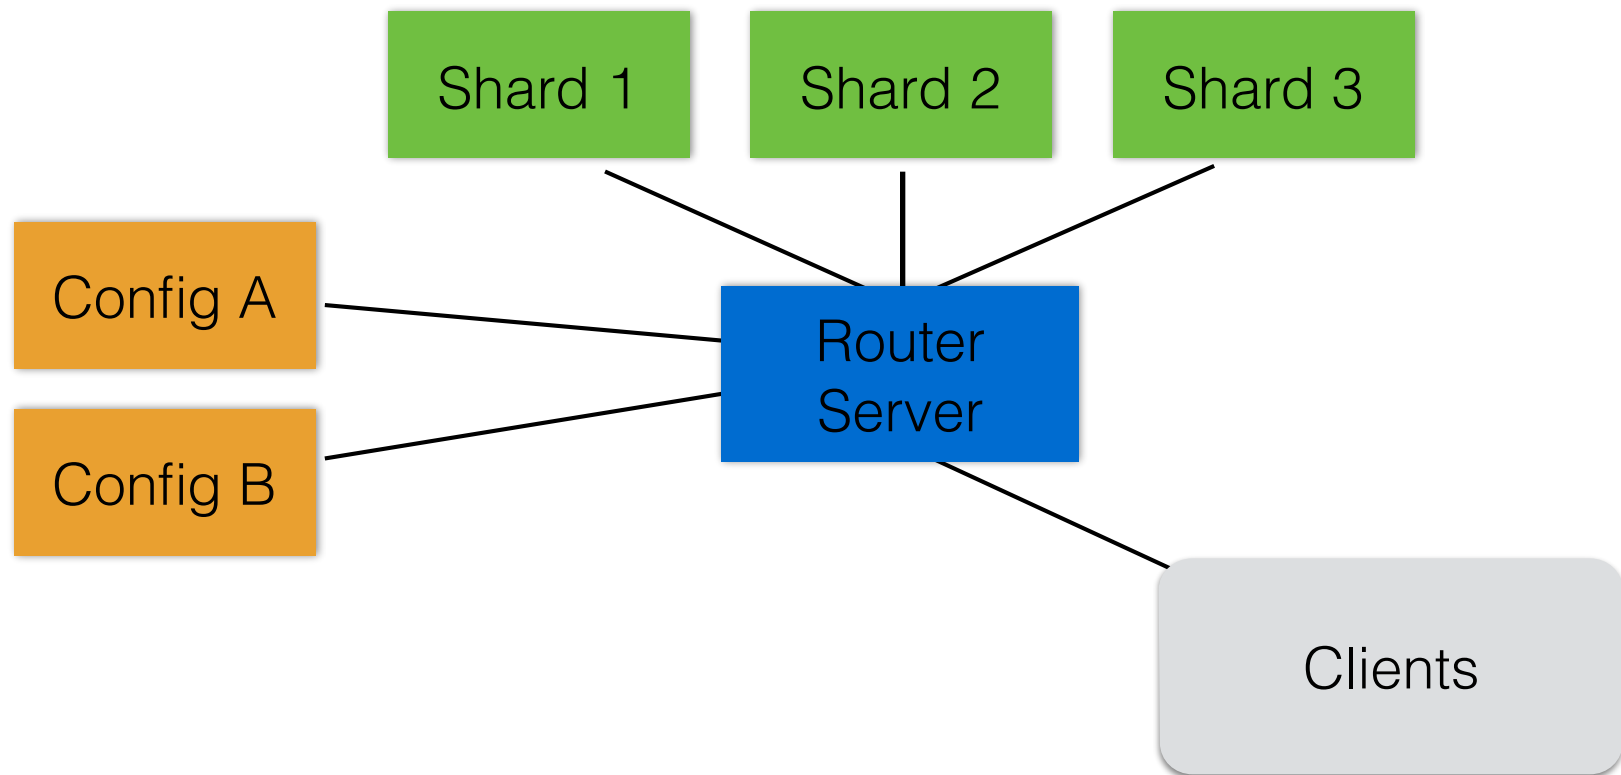
Method	Type	Time (s)	Number of Requests	Time per Request (s)
condor	t2micro	414	1000	0.414
iterative	t2micro	544	1000	0.544
condor	t2small	412	1000	0.412
iterative	t2small	507	1000	0.507
condor	t2medium	409	1000	0.409
iterative	t2medium	491	1000	0.491
condor	m3medium	408	1000	0.408
iterative	m3medium	490	1000	0.490
condor	m3large	402	1000	0.402
iterative	m3large	483	1000	0.483
condor	m3xl	394	1000	0.394
iterative	m3xl	481	1000	0.481

Vertical Scaling Results



Horizontal Scaling

- Sharding-Splitting collection across servers based on content



Horizontal Scaling Results

Testing In Progress

Questions?