

# Load Test Report

**Date:** 6/9/2015

**Test from :** virginia

**Query URL:** http://cloud.kevinohashibenchmark.com/

**Started at:** Tue Jun 9 2015, 01:15:28 -04:00

**Finished at:** Tue Jun 9 2015, 01:16:28 -04:00

**Test link:** https://www.blitz.io/to#/play/input/virginia:15a02c9373c598e9059056ba8a216e47

## Analysis

This rush generated **35,721** successful hits in **60 seconds** and we transferred **486.41 MB** of data in and out of your app. The average hit rate of **595/second** translates to about **51,438,240** hits/day.

The average response time was **238 ms**.

You've got bigger problems, though: **11.07%** of the users during this **rush** experienced timeouts or errors!

Hits **88.93%** (35721)

Errors **0.19%** (75)

Timeouts **10.88%** (4371)

Response Times	Test Configuration	Other Stats
Fastest: <b>82</b> ms	Region: <b>virginia</b>	Avg. Hits: <b>595</b> /sec
Slowest: <b>491</b> ms	Duration: <b>60</b> seconds	Transferred: <b>5.90</b> MB
Average: <b>238</b> ms	Load: <b>1-2000</b> users	Received: <b>480.50</b> MB

## Hits

This rush generated **35,721** successful hits. The number of hits includes all the responses listed below. For example, if you only want **HTTP 200 OK** responses to count as Hits, then you can specify **--status 200** in your rush.

Code	Type	Description	Amount	HITS	HTTP 200 OK <b>100%</b> (35721)
200	HTTP	OK	35721		

## Errors

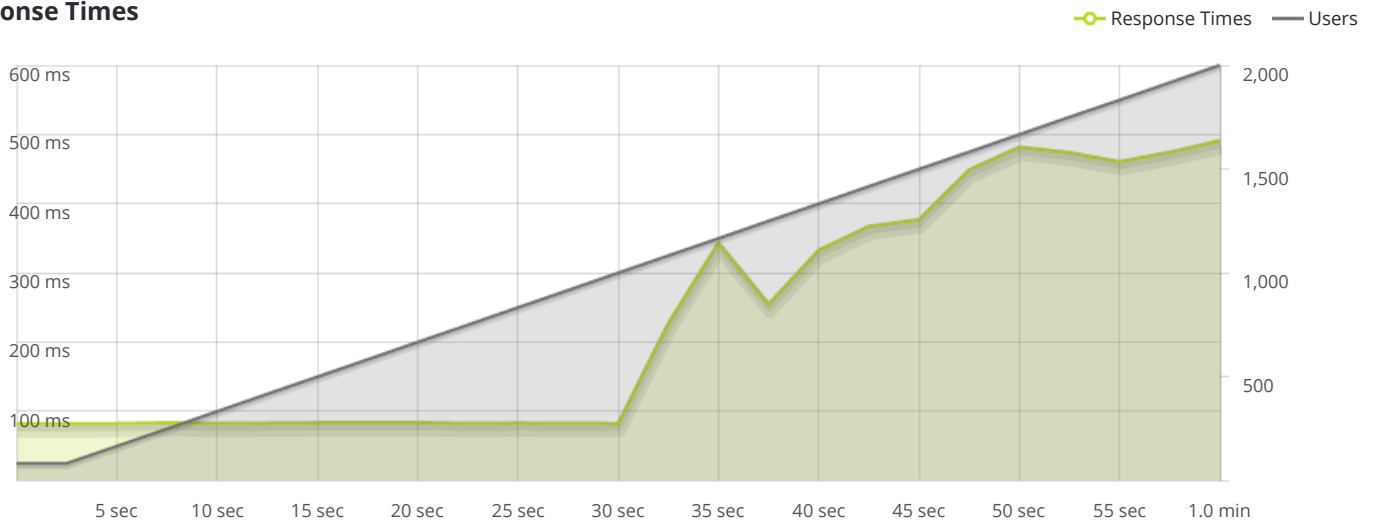
The first error happened at **10 seconds** into the test when the number of concurrent users was at **331**. Errors are usually caused by resource exhaustion issues, like running out of file descriptors or the connection pool size being too small (for SQL databases).

Code	Type	Description	Amount	ERRORS	Connection time... <b>100%</b> (75)
23	TCP	Connection timeout	75		

## Timeouts

The first timeout happened at **32.5 seconds** into the test when the number of concurrent users was at **1082**. Looks like you've been rushing with a timeout of **1000 ms**. Timeouts tend to increase with concurrency if you have lock contention of sorts. You might want to think about in-memory caching using [redis](#), [memcached](#) or [varnish](#) to return stale data for a period of time and asynchronously refresh this data.

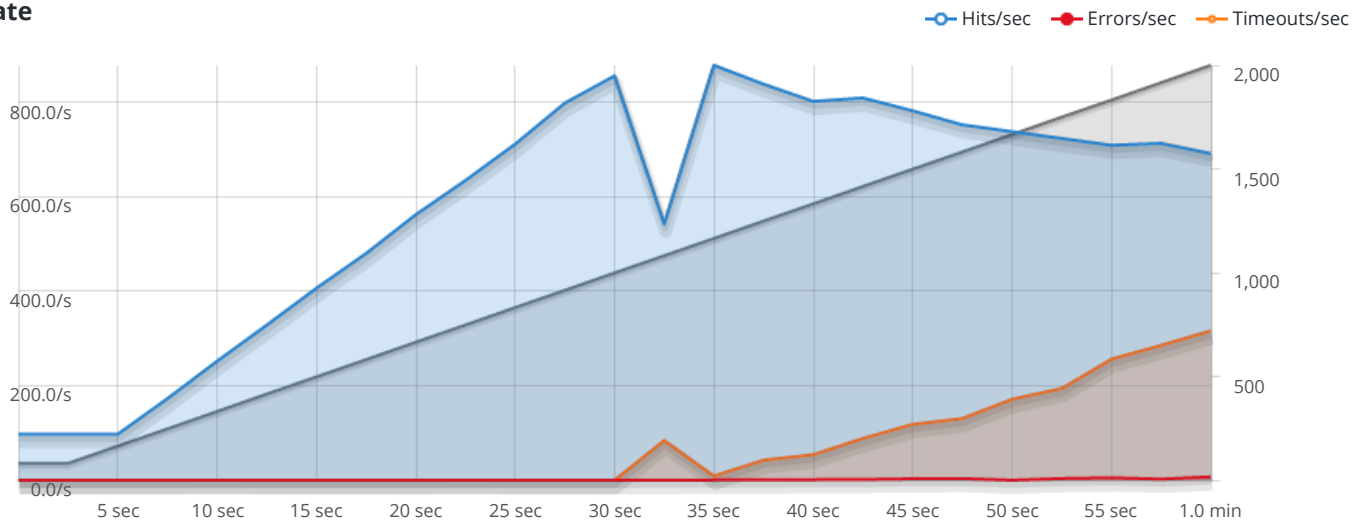
### Response Times



STEP 1  
Response Times

The max response time was: **490 ms @ 2000 users**

### Hit Rate



STEP 1  
Hits/sec Errors/sec Timeouts/sec

The max hit rate was: **878 hits per second**