WHITEPAPER

Pulsar Efficacy Experiment: Testing of NS1 RUM Routing.

NS1. + mapbox

+1.855.GET.NSONE (6766) · SALES@NS1.COM

Table of Contents

List of Figures

1. Executive Summary

An experiment was designed and executed to test the efficacy of Pulsar vs round robin (Shuffle) routing between four major CDNs (Akamai, Fastly, Cloudflare, and Highwinds). The third-party monitoring company Catchpoint was used to simultaneously test both a Pulsar and a Shuffle routed domain from a globally distributed set of test nodes.

This experiment showed that the Pulsar enabled domain exhibited on average 26% lower mean RTT values than the domain that was not using Pulsar routing. Pulsar demonstrated improvements in the mean standard deviation, median, 90th, and 95th percentile round trip times (RTT) across all five nodes tested.

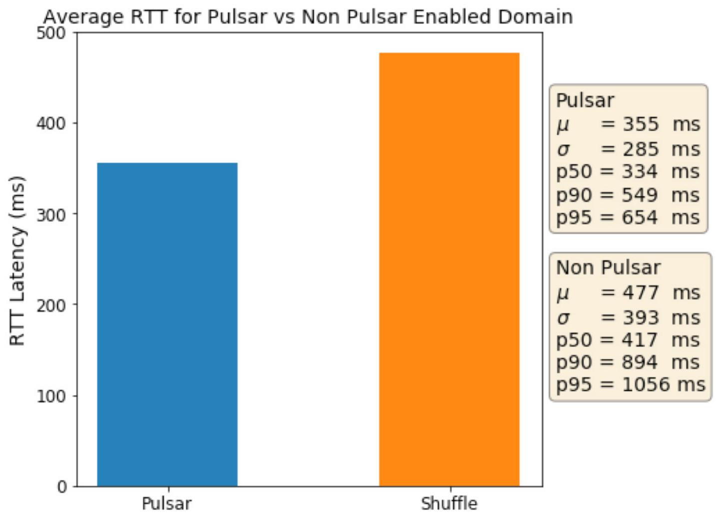

2. Introduction

2.1. Objective

This experiment seeks to quantify the value of using Real User Measurements (RUM) and Pulsar to route end users to a multi-CDN enabled website. Specifically, we would like to optimize for the roundtrip time of a user's request to a given webpage.

2.2. The Merits of Synthetic Monitoring

The decision to use synthetic monitoring was done as it can provide a very highly controlled group of timed web requests from a geographically diverse set of endpoints. We understand that synthetic monitoring has many pitfalls such as: missing most networks (nodes only in a couple of networks), missing most geographies (nodes contained in a couple of geographies), and not testing page loads under real user conditions [1]. However, just because it does not paint a full picture it does not mean that you cannot draw any conclusions from synthetic monitoring. It is very common in statistics to make estimates about the properties of a population given a random sample of the population, this is known as statistical inference.

3. Experimental Setup

3.1. Summary

The experiment is setup with two domains: pulsar.frazao.ca and not-pulsar.frazao.ca which direct traffic using Pulsar and round robin DNS (Shuffle) respectively. When a request is made for either domain one of four CDN endpoints (Fastly, Akamai, Cloudflare, and Highwinds) are chosen depending on the DNS routing. Catchpoint is used to monitor these two domains from five globally distributed nodes for 48 hours.

These four CDNs were chosen as they are recognized as all being high quality CDNs that a company could conceivably use together in some configuration. Some current NS1 customers do route to CDN endpoints using simple some sort of Shuffle Filter since it is often not obvious how to route user traffic to a globally anycasted CDN.

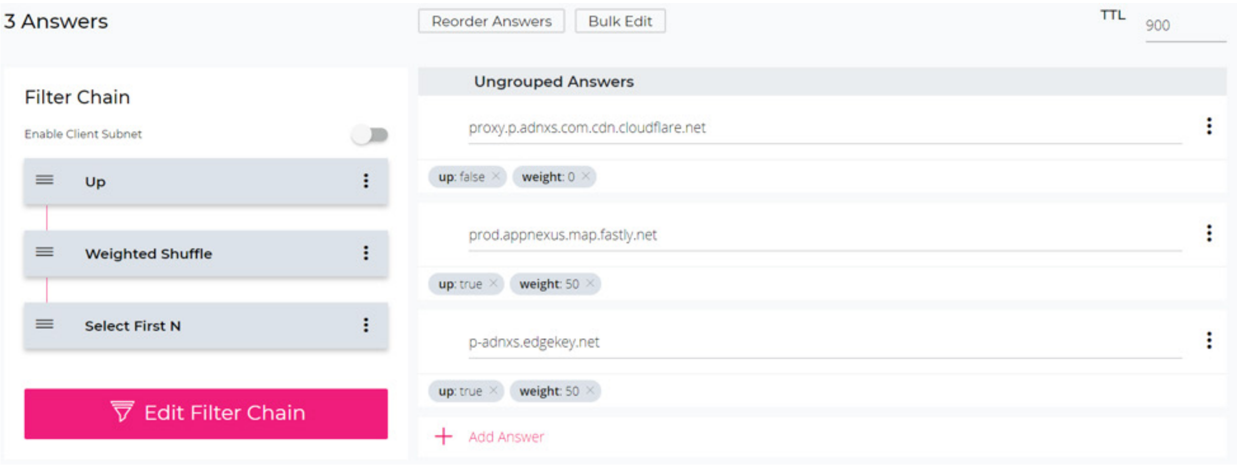

FIGURE 1: EXAMPLE RECORD (1M Q/DAY) USING SHUFFLE TO ROUTE CDNS

3.2. DNS Setup

For this experiment we have configured two domains: pulsar.frazao.ca (Pulsar routed – the experimental group) and not-pulsar.frazao.ca (not Pulsar routed – the control group). These domains could route the user to retrieve a 1x1 pixel at one of four CDNs: Akamai, Fastly, Cloudflare, or Highwinds. The Pulsar enabled record utilized Pulsar for the routing while the non-Pulsar enabled record used round robin DNS (NS1's Shuffle Filter) to randomly choose between each of the 4 CDNs.

Two CNAME records were configured, pulsar.frazao.ca - CNAME and not-pulsar.frazao.ca – CNAME each with the same four answers: fastly.frazao.ca, akamai.frazao.ca, cloudflare.frazao.ca, and highwinds.frazao.ca. Each one of these four answers subsequently gave the user the IP address for a different EC2 server. This setup is demonstrated in the figure below.

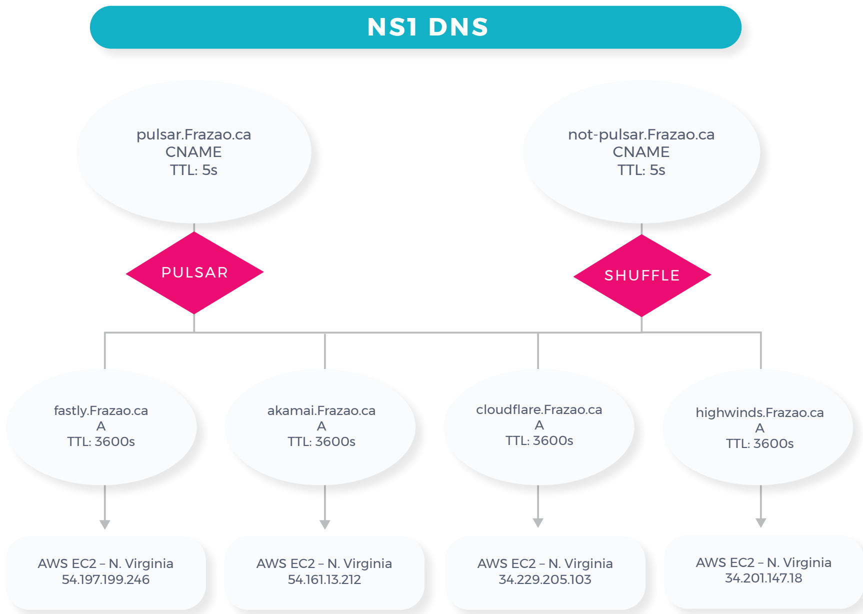

FIGURE 2: EXPERIMENT DNS SETUP

3.3. Webserver Setup

Four webservers were setup on four separate AWS EC2 Instances. Each webserver was configured to accept the host header of either pulsar.frazao.ca or not-pulsar.frazao.ca and to subsequently redirect the requester to one of the four CDNs (e.g., fastly.frazao.ca will redirect to the pixel hosted on Fastly).

Apache was used as the webservers on these EC2 instances and the redirects were accomplished via a line in the virtual host conf file.

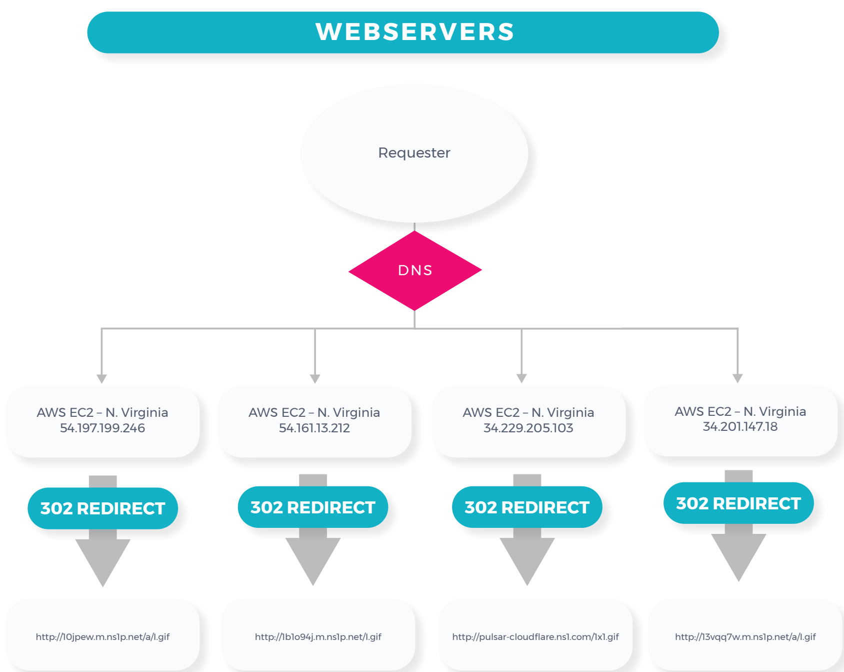

FIGURE 3: EXPERIMENT WEBSERVER SETUP

```
<VirtualHost *:80>
        # The ServerName directive sets the request scheme, hostname and port that
        # the server uses to identify itself. This is used when creating
        # redirection URLs. In the context of virtual hosts, the ServerName
        # specifies what hostname must appear in the request's Host: header to
        # match this virtual host. For the default virtual host (this file) this
        # value is not decisive as it is used as a last resort host regardless.
        # However, you must set it for any further virtual host explicitly.
        #ServerName www.example.com

        ServerAdmin webmaster@frazao.ca
        ServerName pulsar.frazao.ca
        ServerAlias not-pulsar.frazao.ca
        Redirect / http://10jp2ew.m.ns1p.net/a/l.gif
        DocumentRoot /var/www/pulsar.frazao.ca/public_html

        # Available loglevels: trace8, ..., trace1, debug, info, notice, warn,
        # error, crit, alert, emerg.
        # It is also possible to configure the loglevel for particular
        # modules, e.g.
        #LogLevel info ssl:warn

        ErrorLog ${APACHE_LOG_DIR}/error.log
        CustomLog ${APACHE_LOG_DIR}/access.log combined

        # For most configuration files from conf-available/, which are
        # enabled or disabled at a global level, it is possible to
        # include a line for only one particular virtual host. For example the
        # following line enables the CGI configuration for this host only
        # after it has been globally disabled with "a2disconf".
        #Include conf-available/serve-cgi-bin.conf
</VirtualHost>

# vim: syntax=apache ts=4 sw=4 sts=4 sr noet
```

FIGURE 4: EXAMPLE VIRTUAL HOST CONF FILE

3.4. Catchpoint Monitoring

The Catchpoint tests setup for this experiment were designed to control for as many variables as possible so that we are only testing the effectiveness of Pulsar and not some other factor that we did not control for. To this end two Catchpoint tests were setup to test both pulsar.frazao.ca and not-pulsar.frazao.ca simultaneously from the same five nodes.

The “Object” monitor type was used, with a five-minute frequency, running the tests concurrently from five nodes (Paris – Cogent, New York – Level3, Johannesburg – Vox, Tokyo – SoftLayer, and Sao Paulo - AWS). These two tests were run for 48 hours.

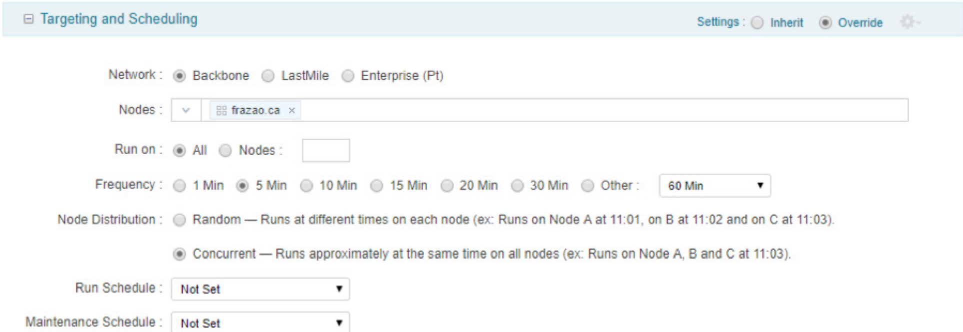

FIGURE 5: TARGETING AND SCHEDULING CATCHPOINT TEST SETTINGS

4. Results & Analysis

Approximately 5,750 runs were conducted by Catchpoint from the five nodes (~2,880 runs per domain). All analysis was conducted using Python and Jupyter Notebooks which is available for review upon request. On average the Pulsar enabled domain had a round trip time of 355ms while the non-Pulsar domain had a round trip time of 477ms. On average the Pulsar domain was 122ms (26%) faster than the domain that did not have Pulsar enabled. The Pulsar enabled domain also exhibited a smaller standard deviation in the RTT and had smaller RTT values at all percentiles tested.

Metric	Pulsar	Non-Pulsar (Shuffle)
Mean	354.91 ms	476.89 ms
Standard Deviation	284.52 ms	393.22 ms
50th Percentile	334.00 ms	417.00 ms
90th Percentile	548.60 ms	894.10 ms
95th Percentile	653.60 ms	1056.1 ms

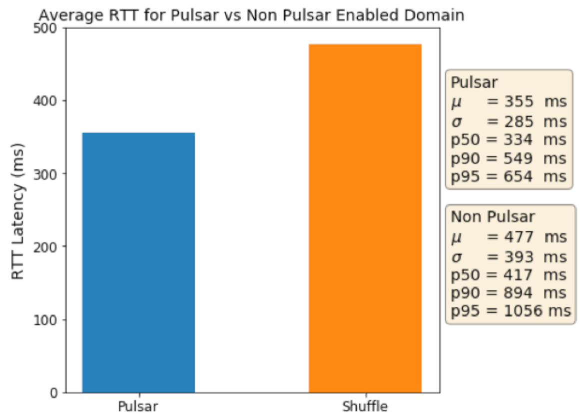

FIGURE 6: AVERAGE RTT, ALL NODES, PULSAR VS SHUFFLE

This observation remained true at every node that was tested where the mean RTT for Pulsar outperformed Shuffle by as much as 181ms (32%) in one instance. In all other statistics examined (standard deviation, median, 90th percentile, and 95th percentile) Pulsar outperformed Shuffle.

	Pulsar					Shuffle				
All stats in ms	μ	σ	p50	p90	p95	μ	σ	p50	p90	p95
Paris FR - Cogent	378	458	250	667	835	559	527	450	944	1055
Johannesburg ZA - Vox	546	128	512	614	741	696	254	545	1070	1091
Tokyo JP - Softlayer	413	124	369	525	625	557	313	461	818	1045
Sao Paulo BR - AWS	319	109	307	350	376	440	262	337	695	719
New York - Level3	118	234	57	242	348	132	290	69	265	358

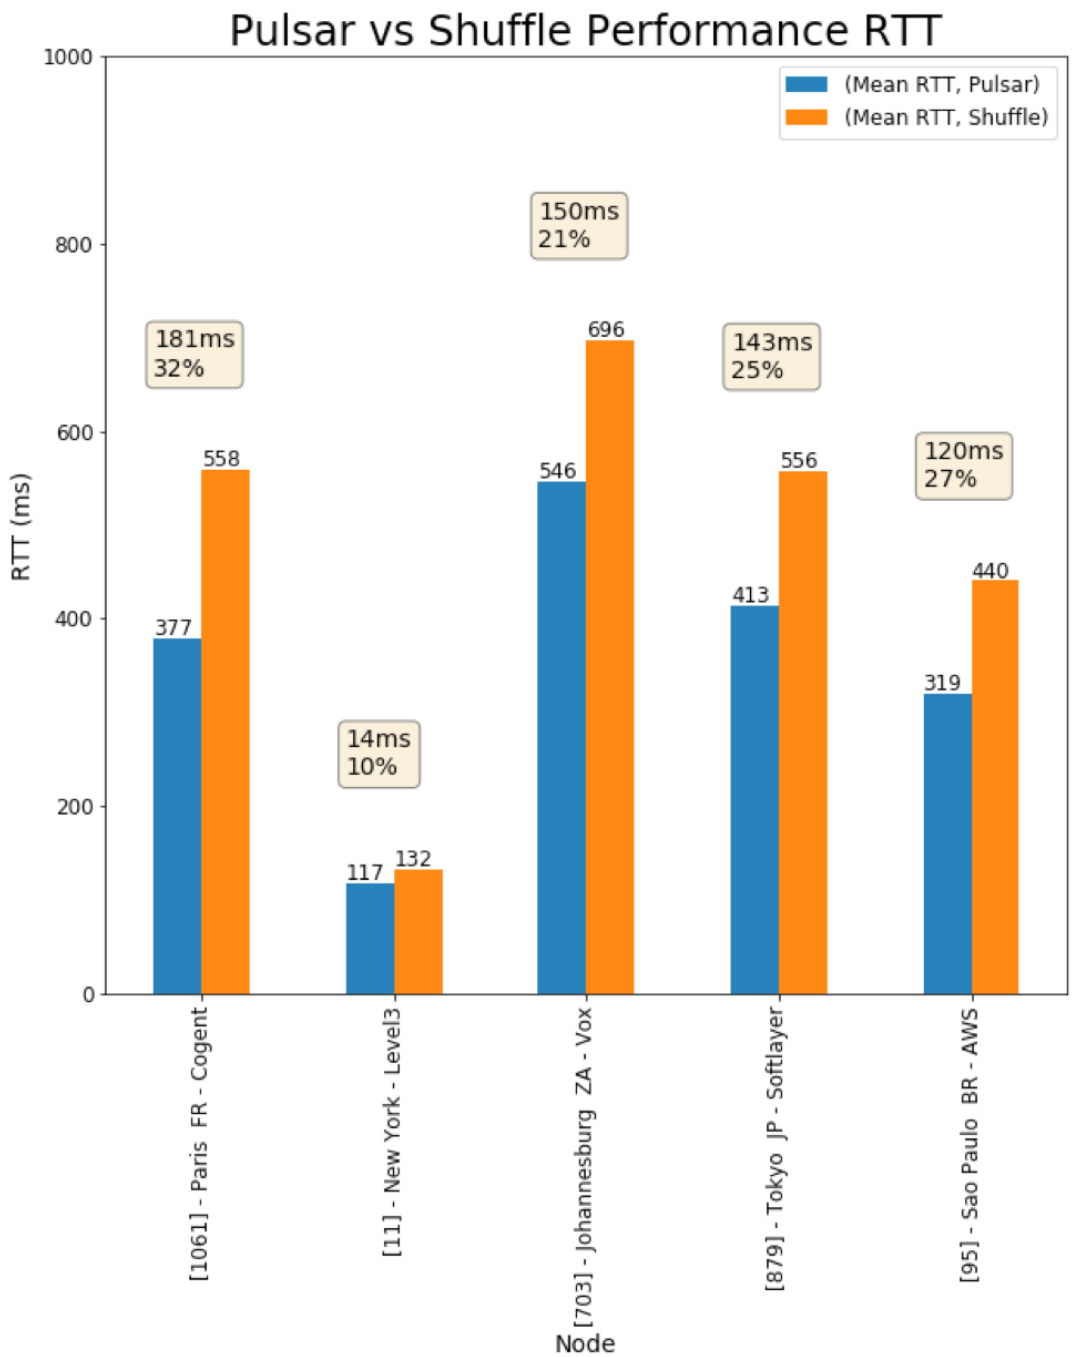

FIGURE 7: AVERAGE RTT BY NODE

5. Next Steps

This experiment's results suggest that Pulsar provides significant advantage over round robin routing for multi-endpoint domains, especially outside of North America. As discussed earlier in the paper, there are some potential problems with using synthetic testing for measuring end user experience, so it would be valuable to reproduce this experiment with many more nodes, over a longer time horizon. Alternatively, we could also try to reproduce this experiment using Catchpoint's "Last Mile" tests to provide better representation of end user networks and see if the results hold.

I believe that the most interesting claim to test would be the efficacy of Pulsar vs geographical routing where either:

1. Certain CDNs are chosen as defaults for a given geographical area – would there be a benefit from using Pulsar in this case (e.g., I have access to Akamai, Cloudflare, and Highwinds, but in Brazil I only use Highwinds)

2. Efficacy of Pulsar versus endpoints with unicast known geographical endpoints (e.g., Is there any reason to use Pulsar to route AWS East 1 vs West 1)

6. Works Cited

[1] P. Mastin, Real User Measurements, B. Anderson, Ed., Sebastopol, California: O'Reilly Media, Inc., 2016.

7. Appendix

7.1. Pulsar vs Shuffle by Node

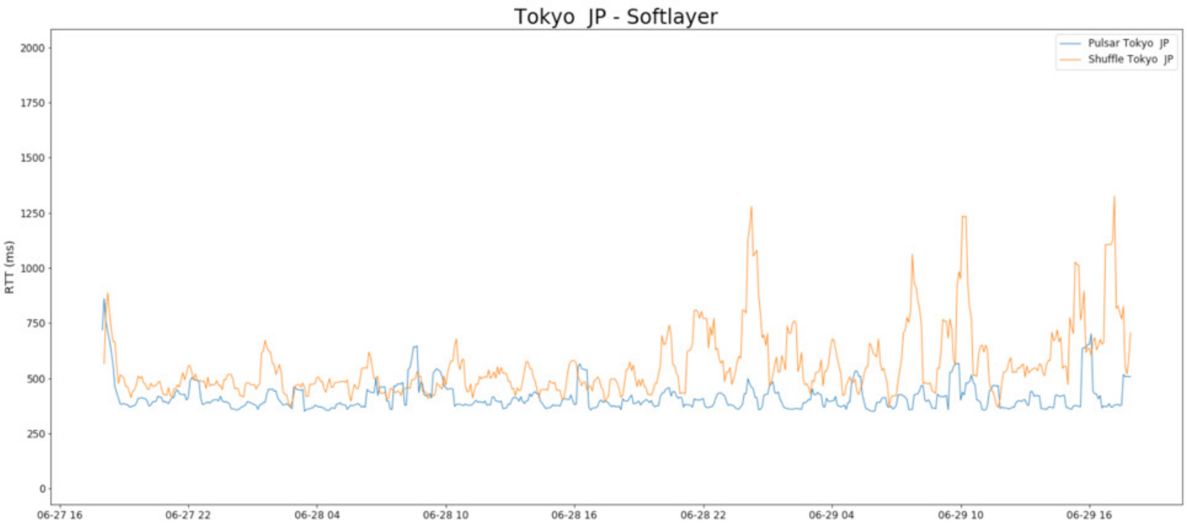

FIGURE 8: PULSAR VS SHUFFLE - TOKYO

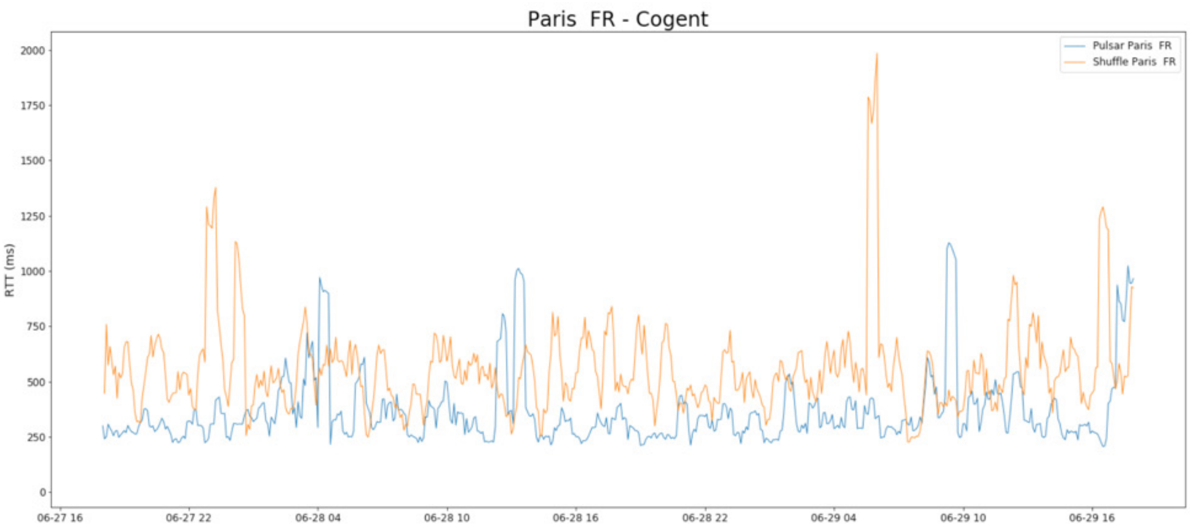

FIGURE 9: PULSAR VS SHUFFLE - PARIS

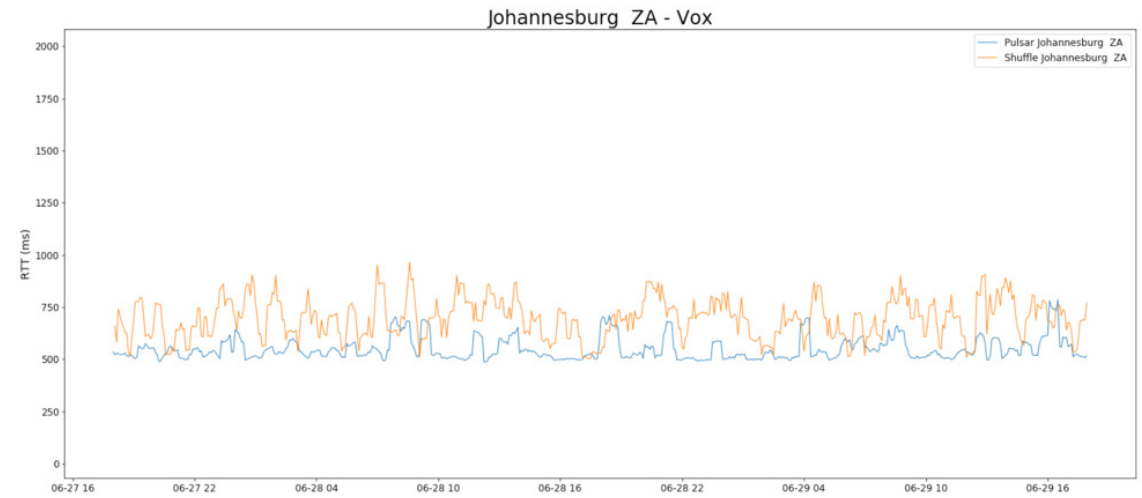

FIGURE 10: PULSAR VS SHUFFLE - JOHANNESBURG

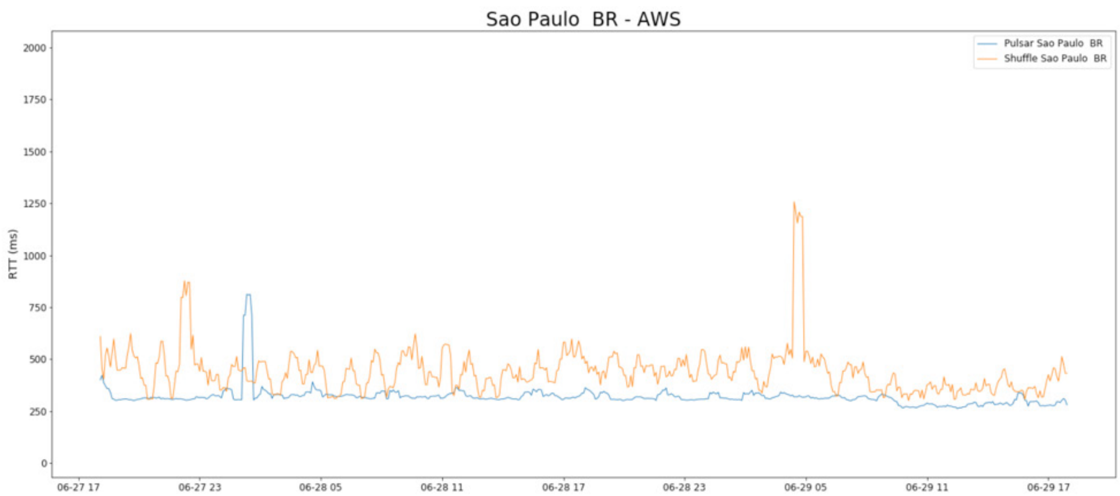

FIGURE 11: PULSAR VS SHUFFLE - SAO PAULO

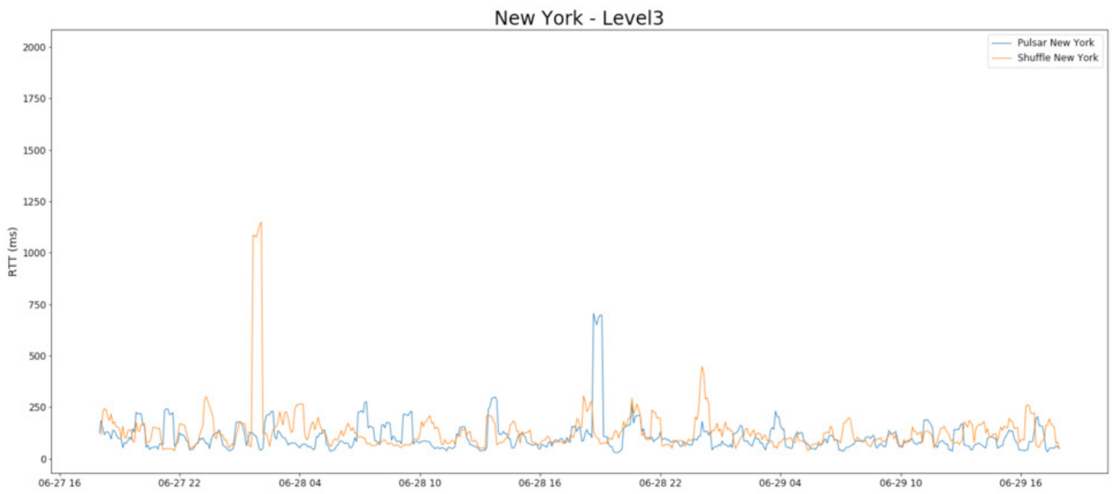

FIGURE 12: PULSAR VS SHUFFLE - NEW YORK

7.2. RTT Difference Between Pulsar and Shuffle by Node and Hour

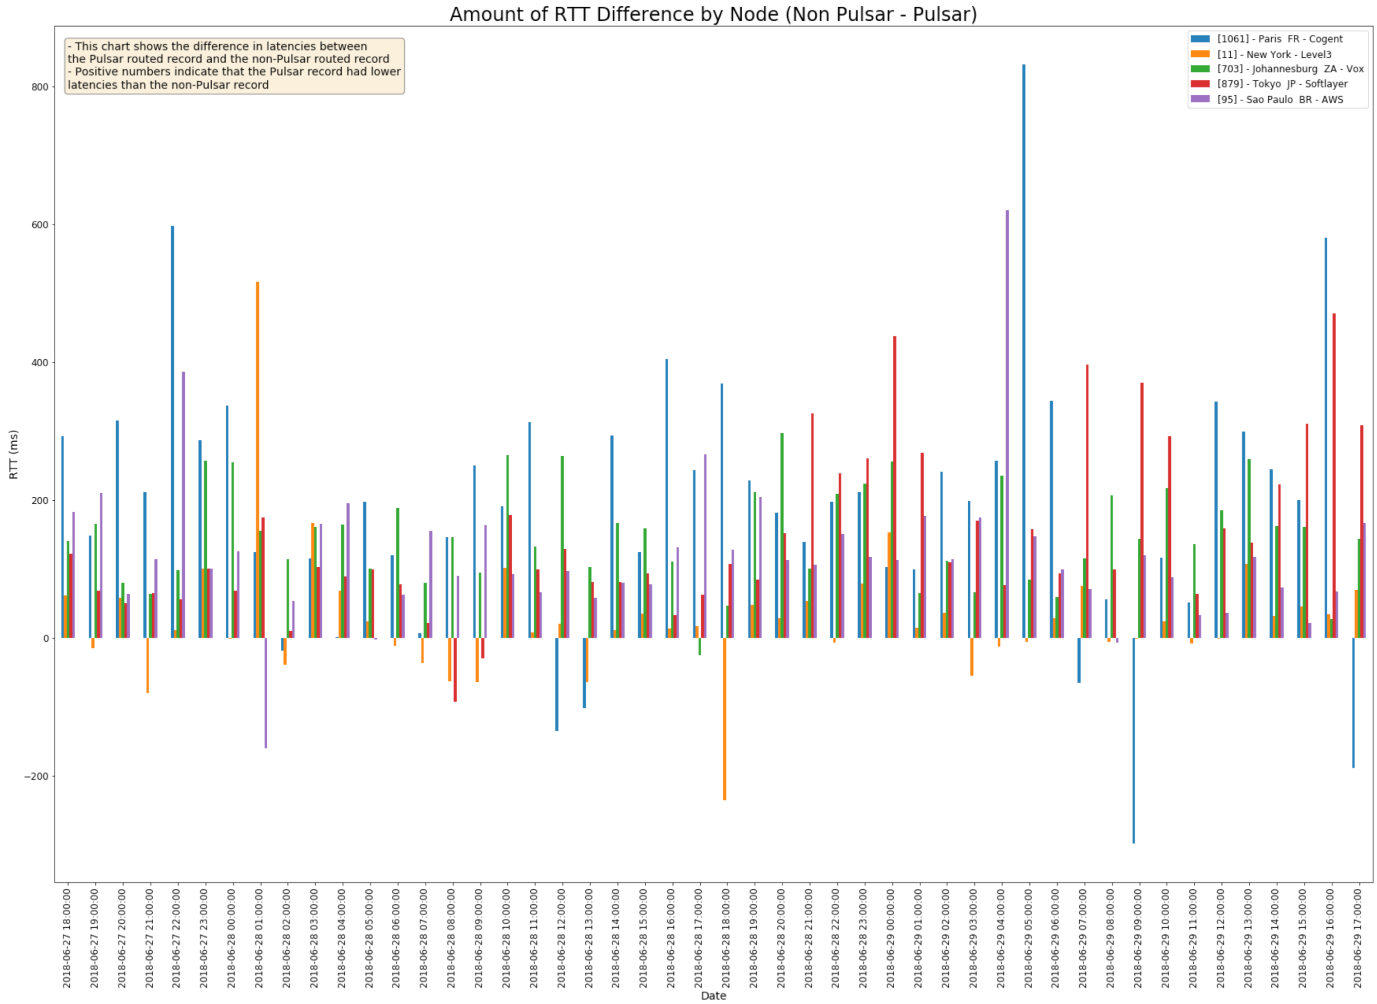

FIGURE 13: RTT DIFFERENCE BY HOUR, PULSAR VS SHUFFLE

7.3. Pulsar vs Shuffle Boxplots

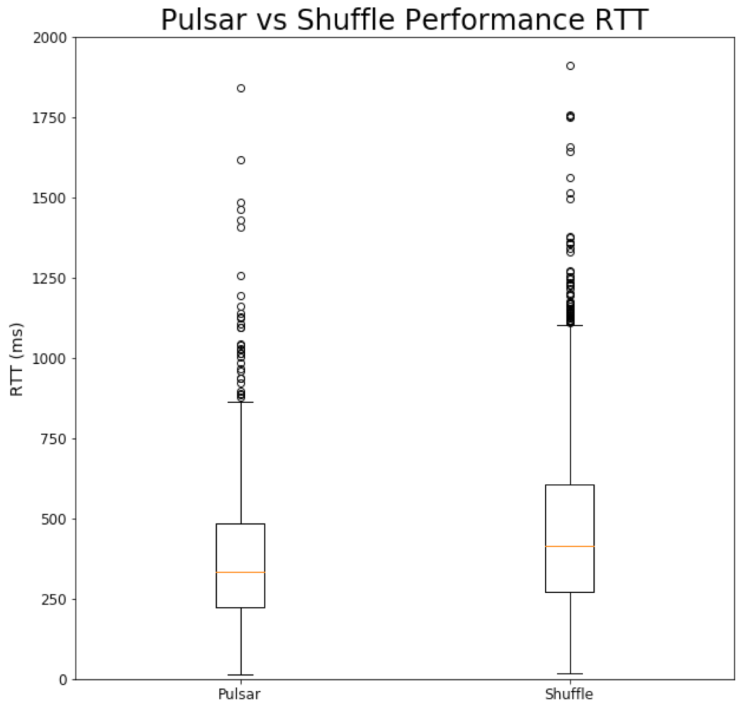

FIGURE 14: PULSAR VS SHUFFLE - BOXPLOT

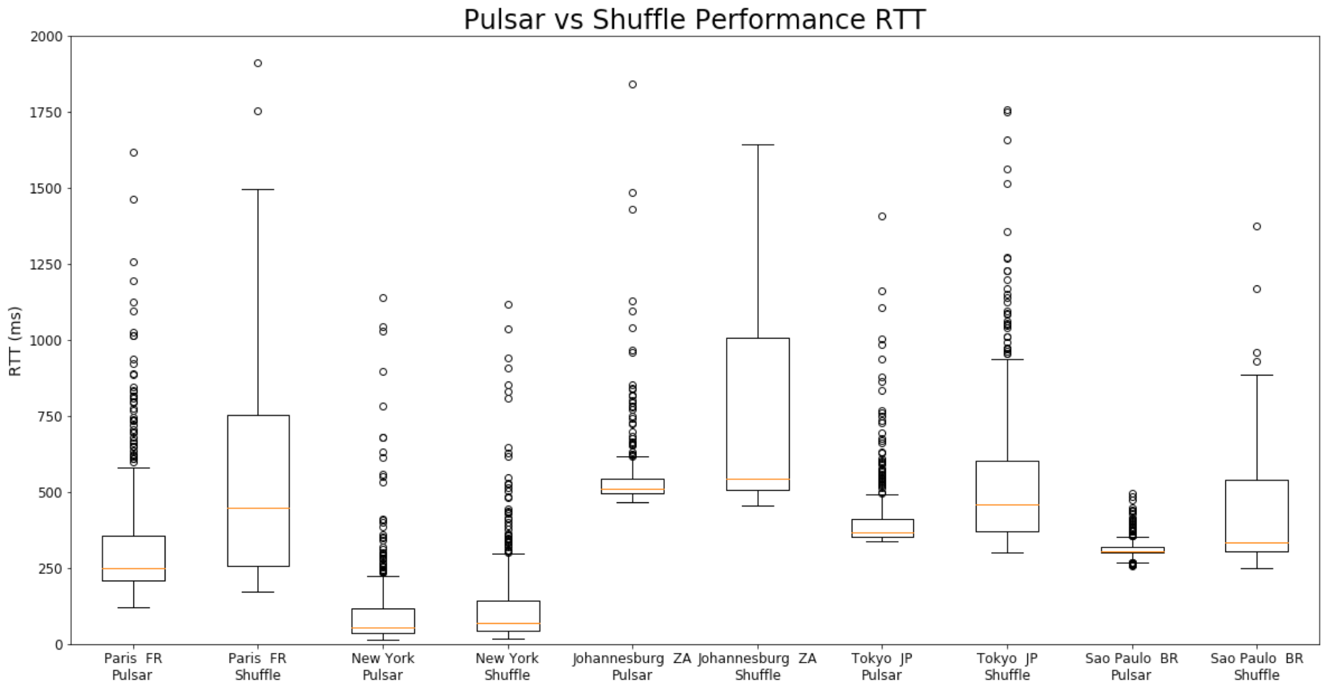

FIGURE 15: PULSAR VS SHUFFLE - BOXPLOT BY NODE

7.4. Cumulative Distribution Function

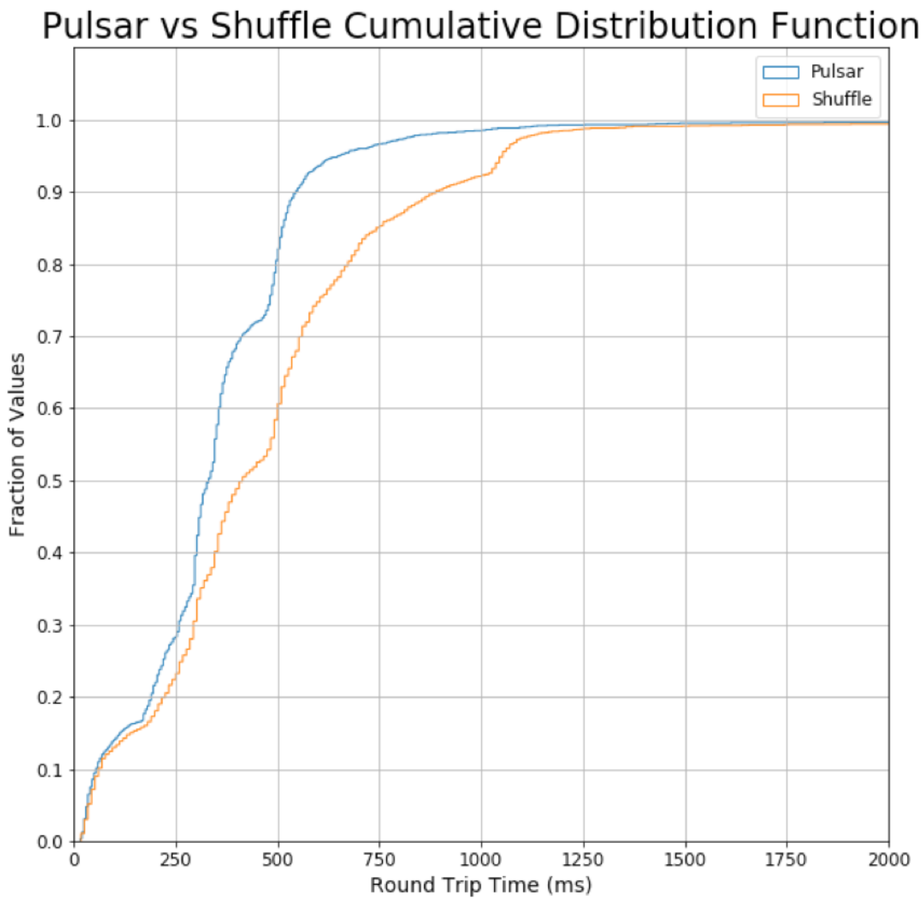

FIGURE 16: PULSAR AND SHUFFLE CDF

ABOUT NS1

NS1 is the leader in next generation DNS solutions that orchestrate the delivery of the world's most critical internet and enterprise applications. Only NS1's purpose-built platform, which is built on a modern API-first architecture, transforms DNS into an intelligent, efficient and automated system, driving dramatic gains in reliability, resiliency, security and performance of application delivery infrastructure. Many of the highest-trafficked sites and largest global enterprises trust NS1, including Salesforce, LinkedIn, Dropbox, Nielsen, Squarespace, Pandora and The Guardian.