Real User Measurement DNS
Using Big Data to decide RData

Samir Jafferali
Staff SRE, LinkedIn

Global load balancing is evolving

- Complex
- Dynamic
- Real-time

Closing thoughts

Questions?

Real User Measurement DNS Using Big Data to decide RData

Samir Jafferali
Staff SRE, LinkedIn

Samir Jafferali
Staff SRE, LinkedIn

Global load balancing is evolving

- Complex
- Dynamic
- Real-time

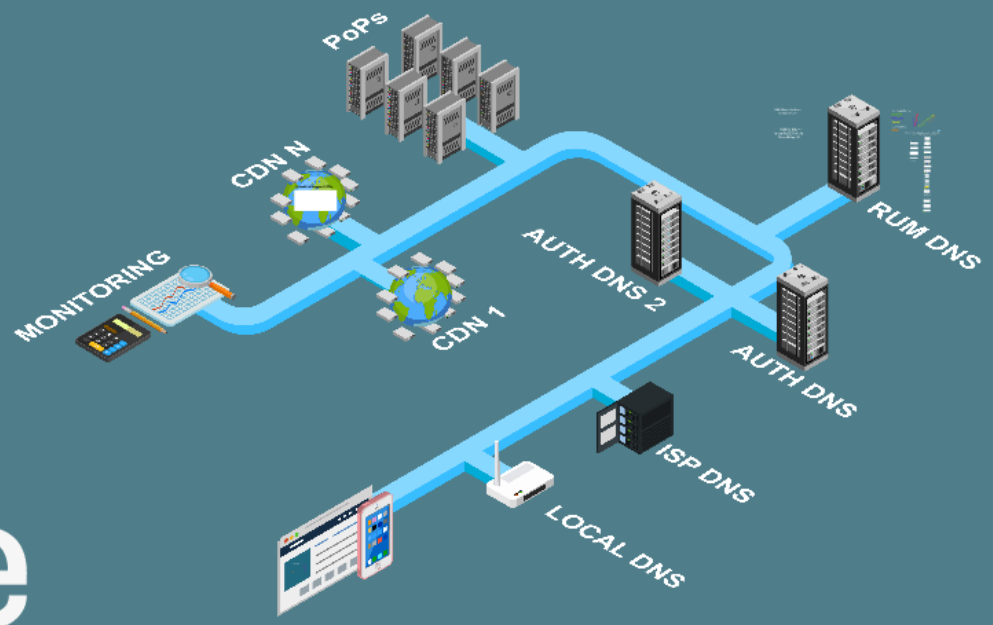

RUM DNS is like Waze
for the Internet.

RUM for DNS
is used by 50% of the
Alexa US top 10

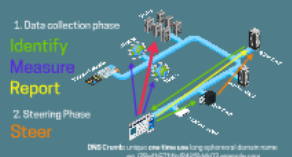

RUM DNS

RUM DNS is like Waze for the Internet.

RUM for DNS (+ Big Data for some) is used by 50% of the Alexa US top 10

1. Data collection phase

Identify

Measure

Report

2. Steering Phase

Steer

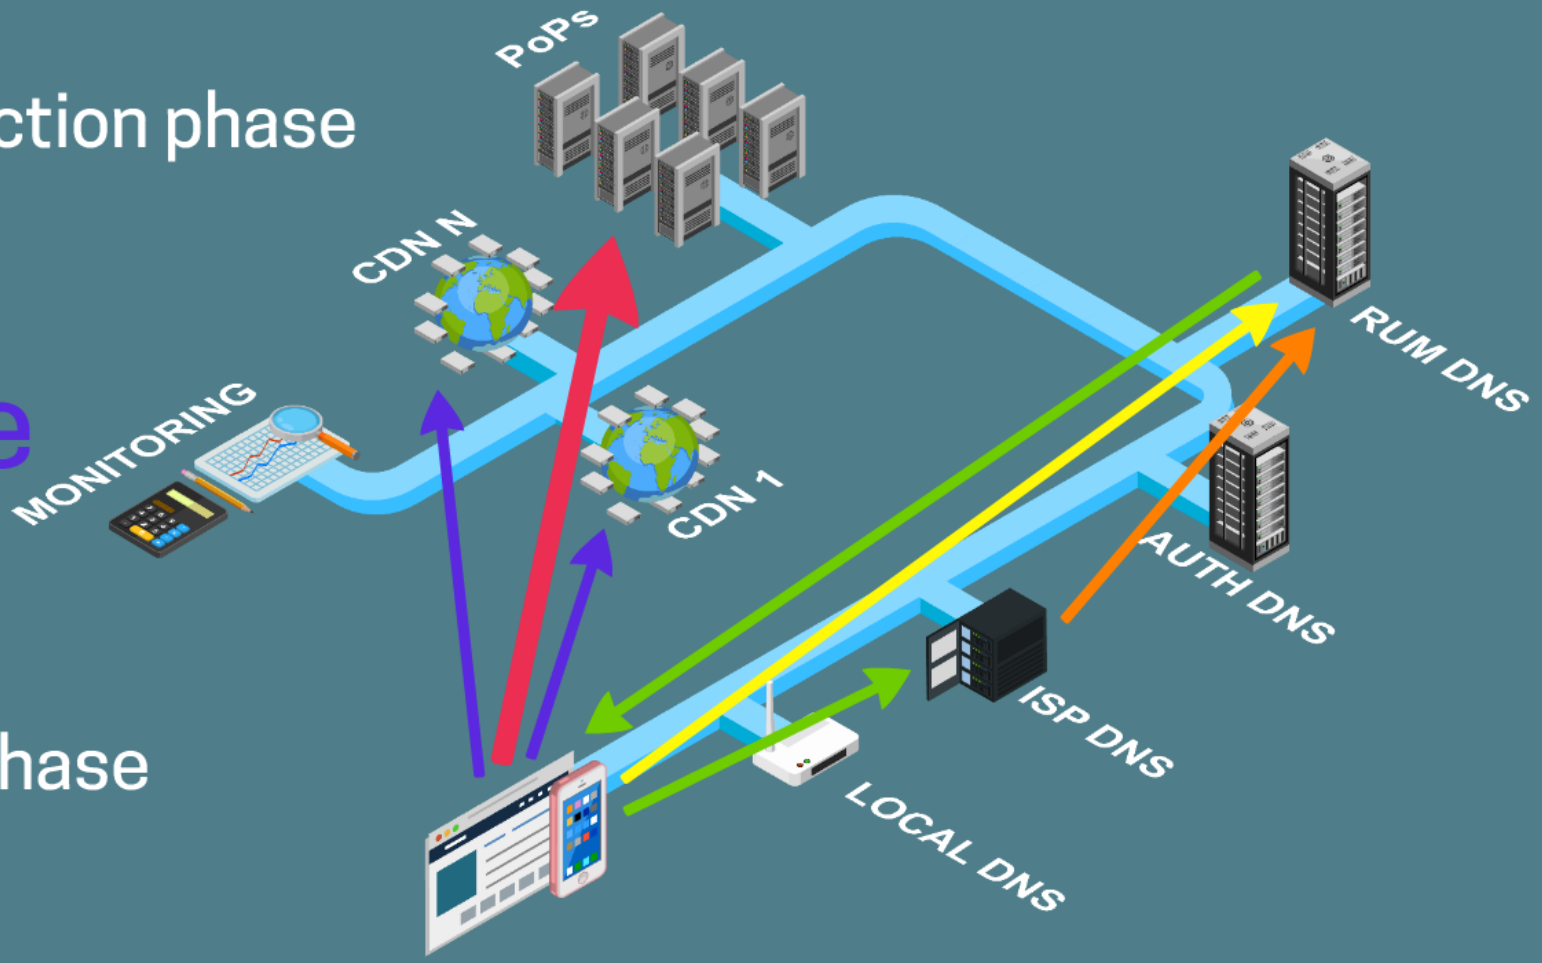

DNS Crumb: unique **one time use** long ephemeral domain name
eg. j38ejfh57fjfkd84jjf8ddk03.example.com

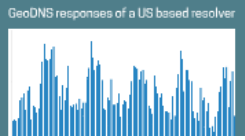

Two CDNs with opposite performance throughout day

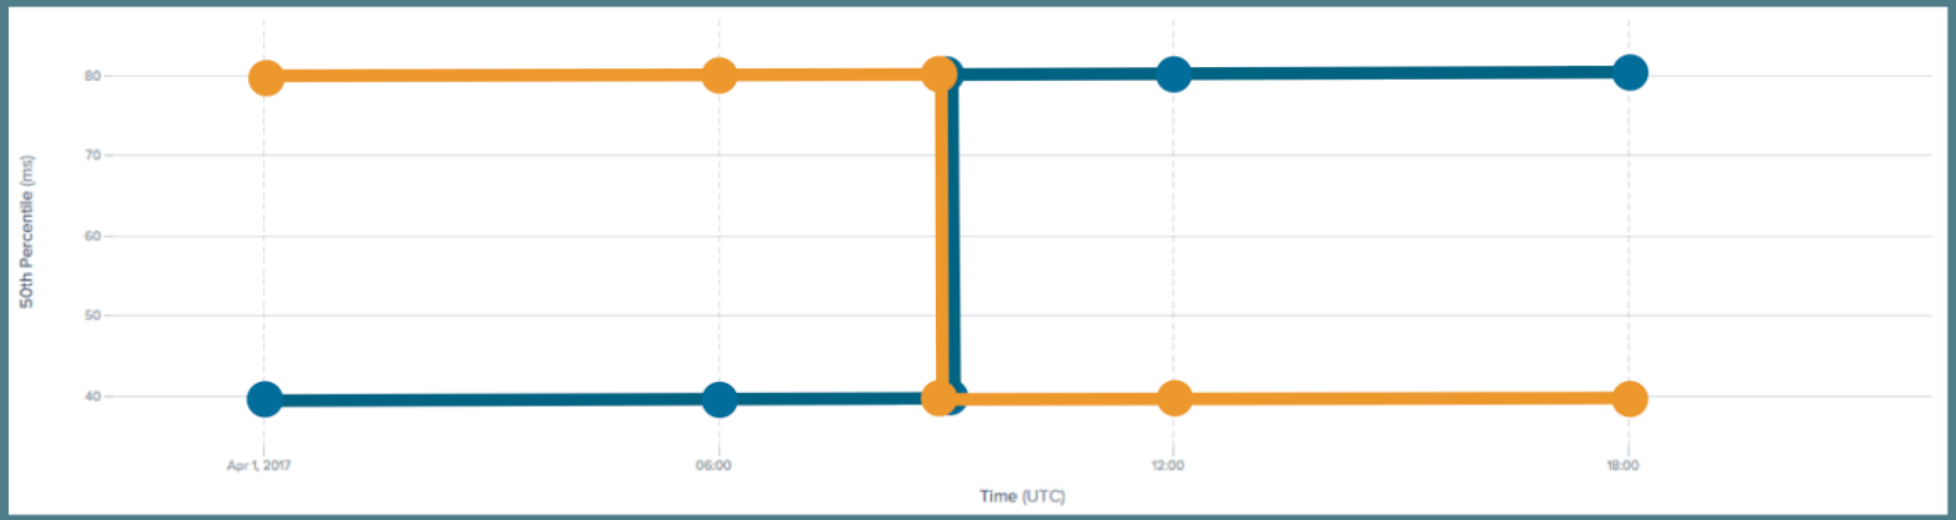

Picking the best CDN per ASN segment

Picking the best CDN per ASN segment improves performance

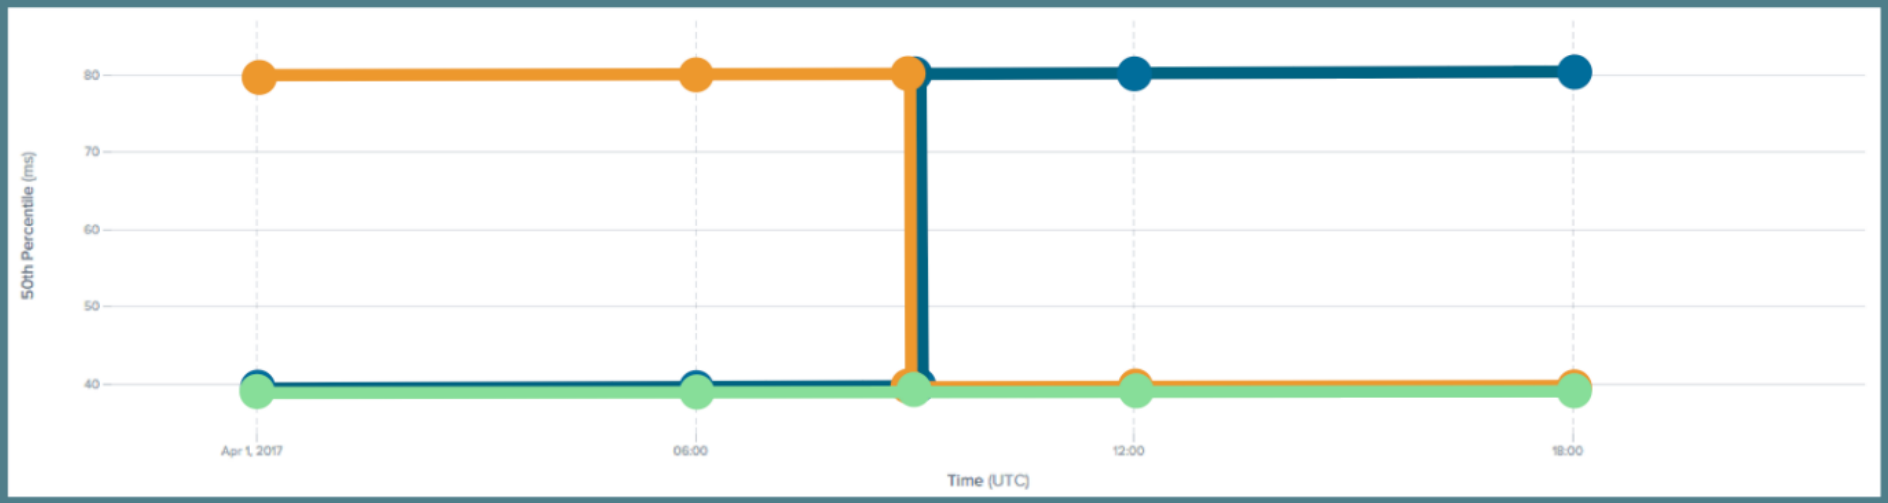

LNKD steered by RUM DNS comprising 3 CDNs

LNKD steered by RUM DNS comprising 3 CDNs

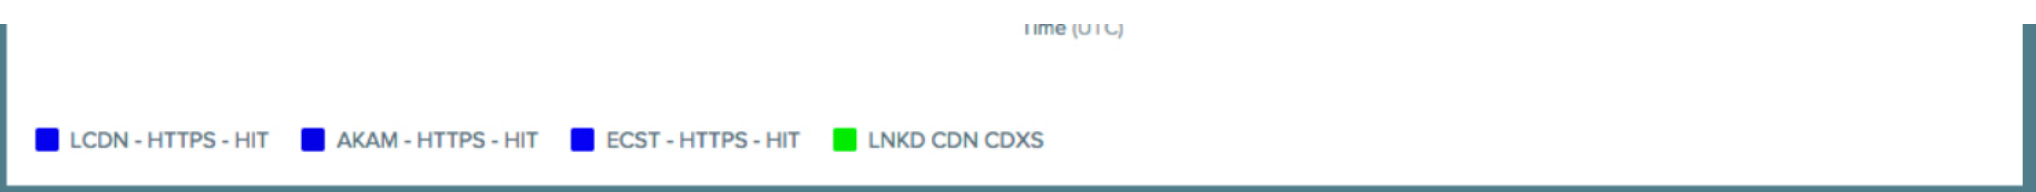

RUM DNS failing out a CDN automatically

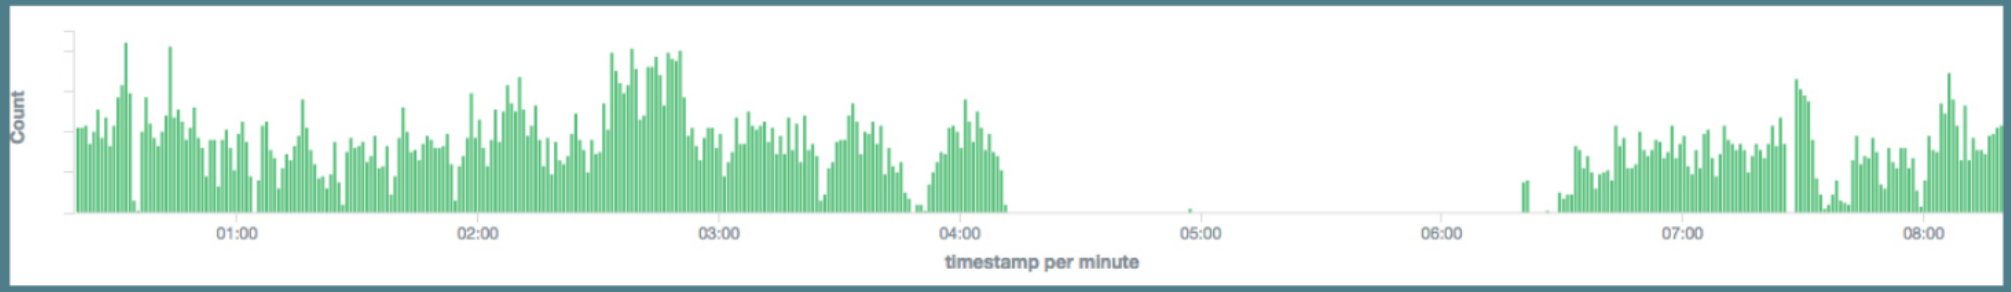

LinkedIn leveraging RUM DNS since 2014
Decided to try PoP steering

1. Data collection phase

Identify

Measure

Report

2. Steering Phase

Steer

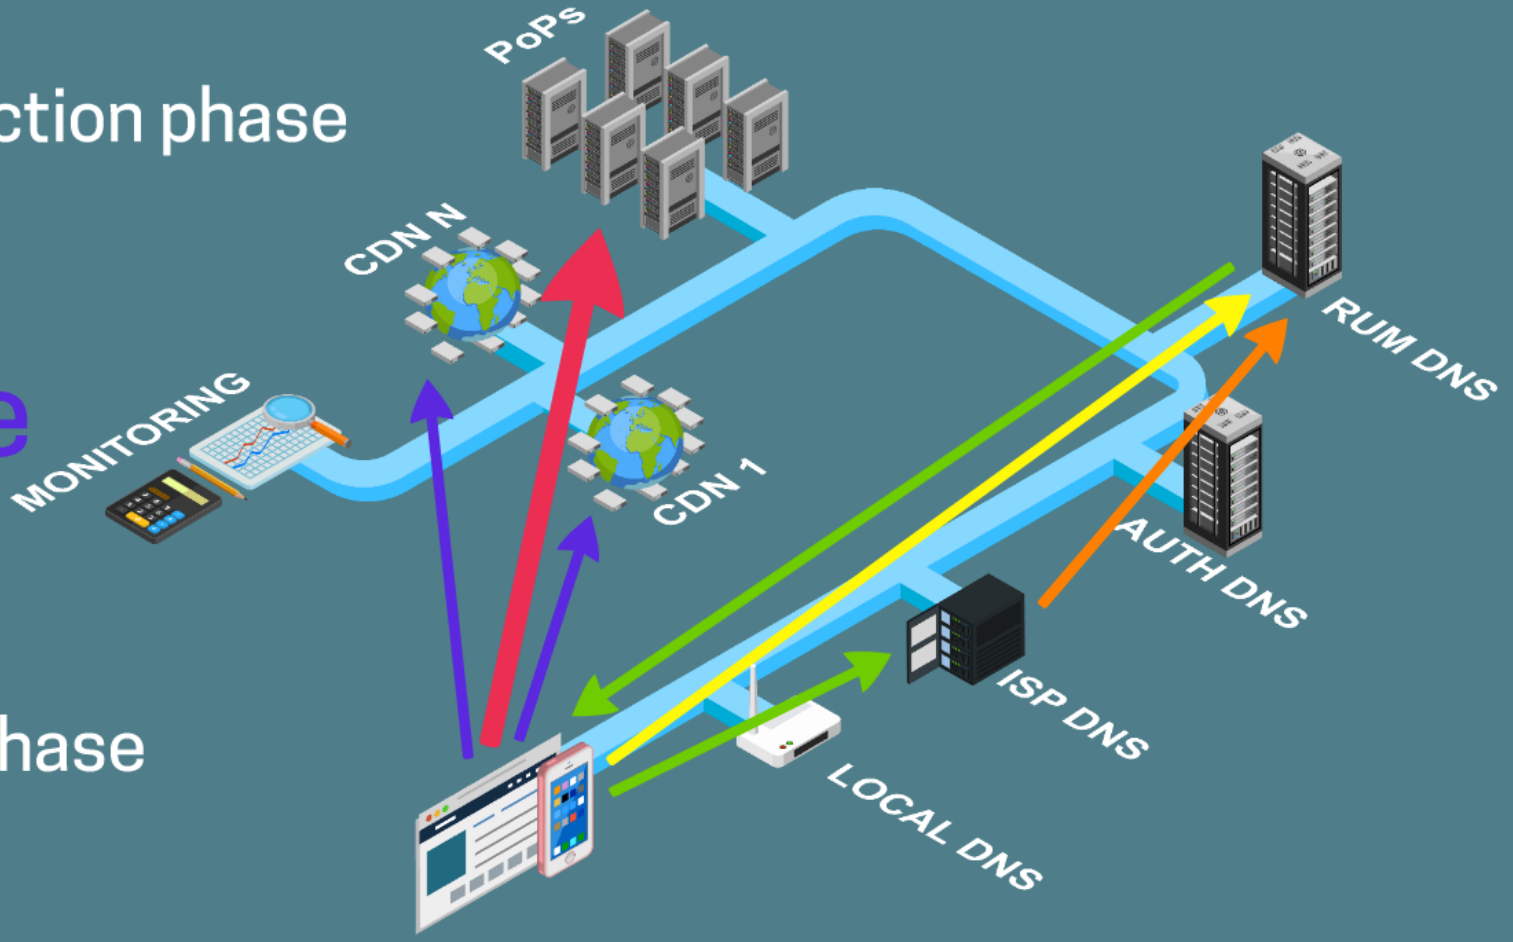

DNS Crumb: unique **one time use** long ephemeral domain name
eg. j38ejfh57fjfkd84jjf8ddk03.example.com

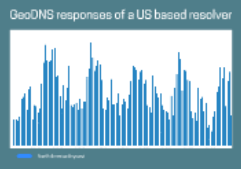

GeoDNS responses of a US based resolver

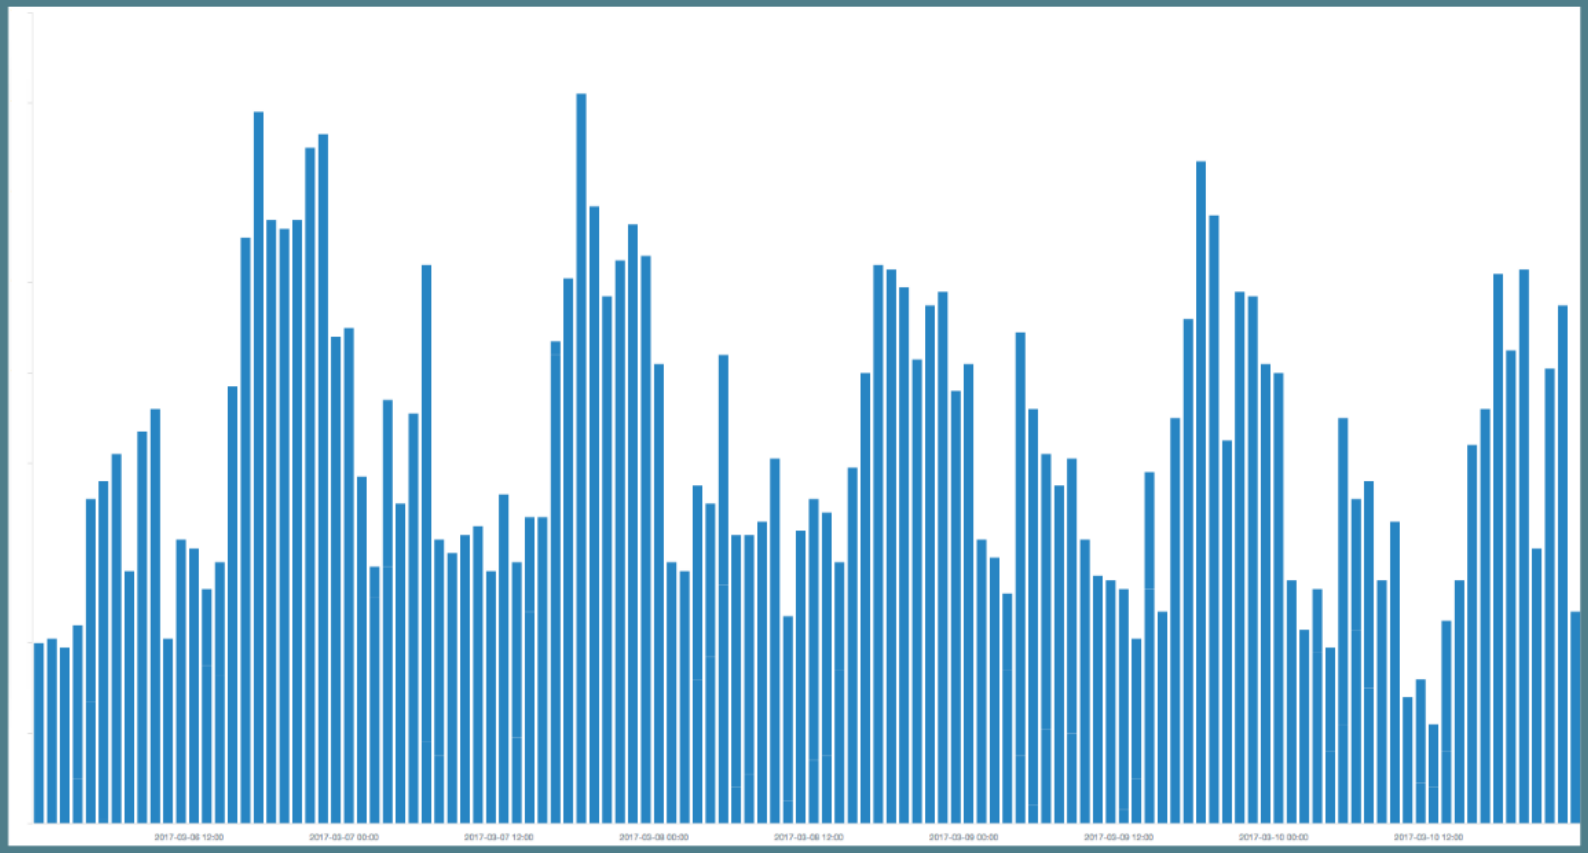

North America Anycast

RUM DNS responses of a US based resolver

North America Anycast

RUM DNS responses of a US based resolver

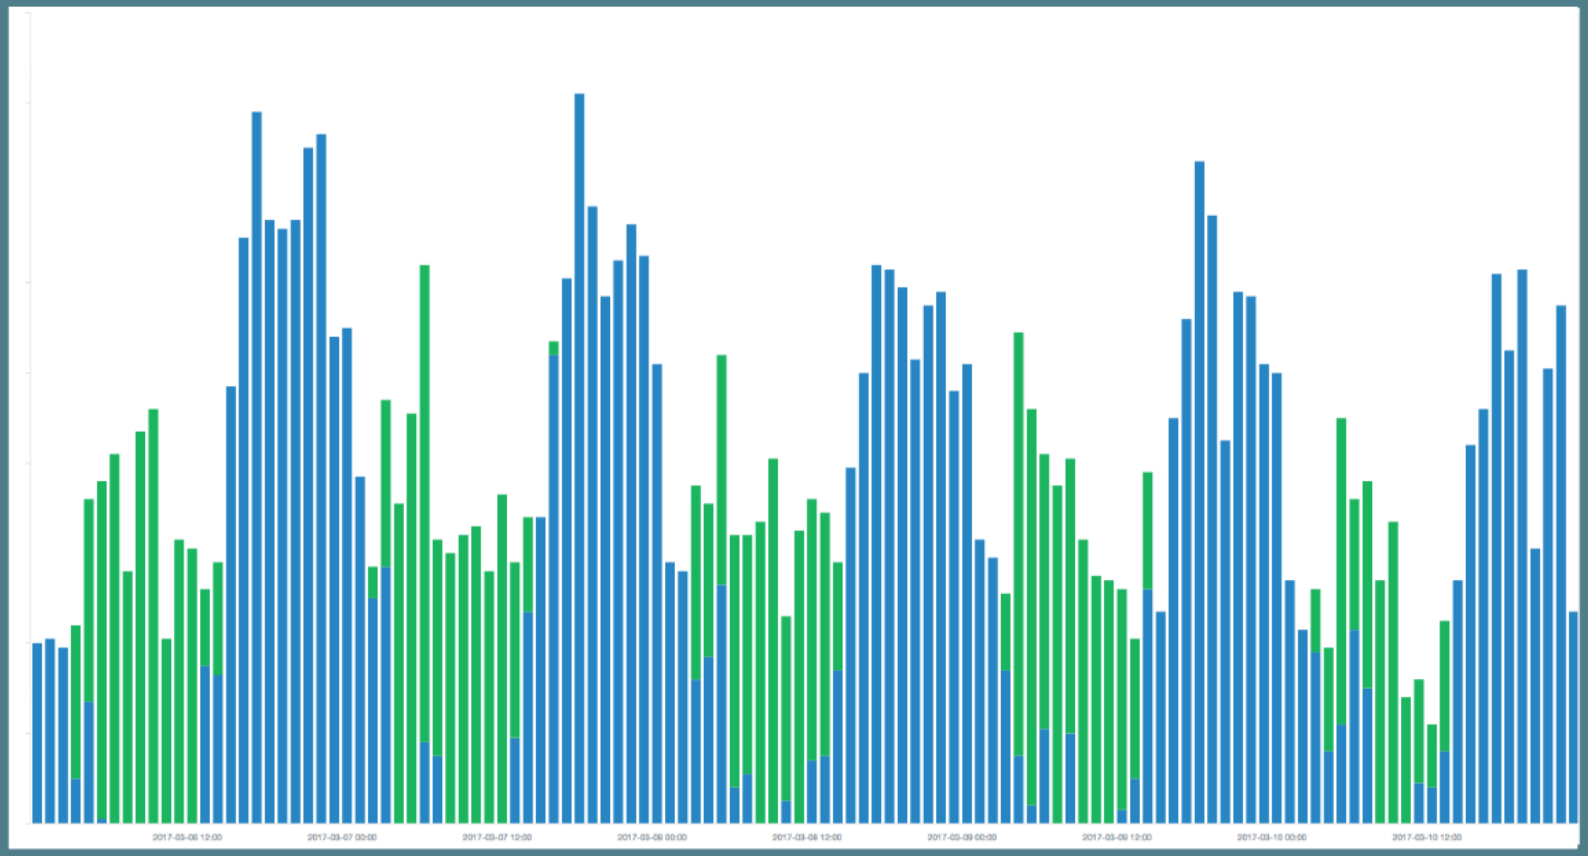

North America Anycast Mumbai PoP

India improves when North America ramped

North America Anycast Mumbai PoP

India improves when North America ramped

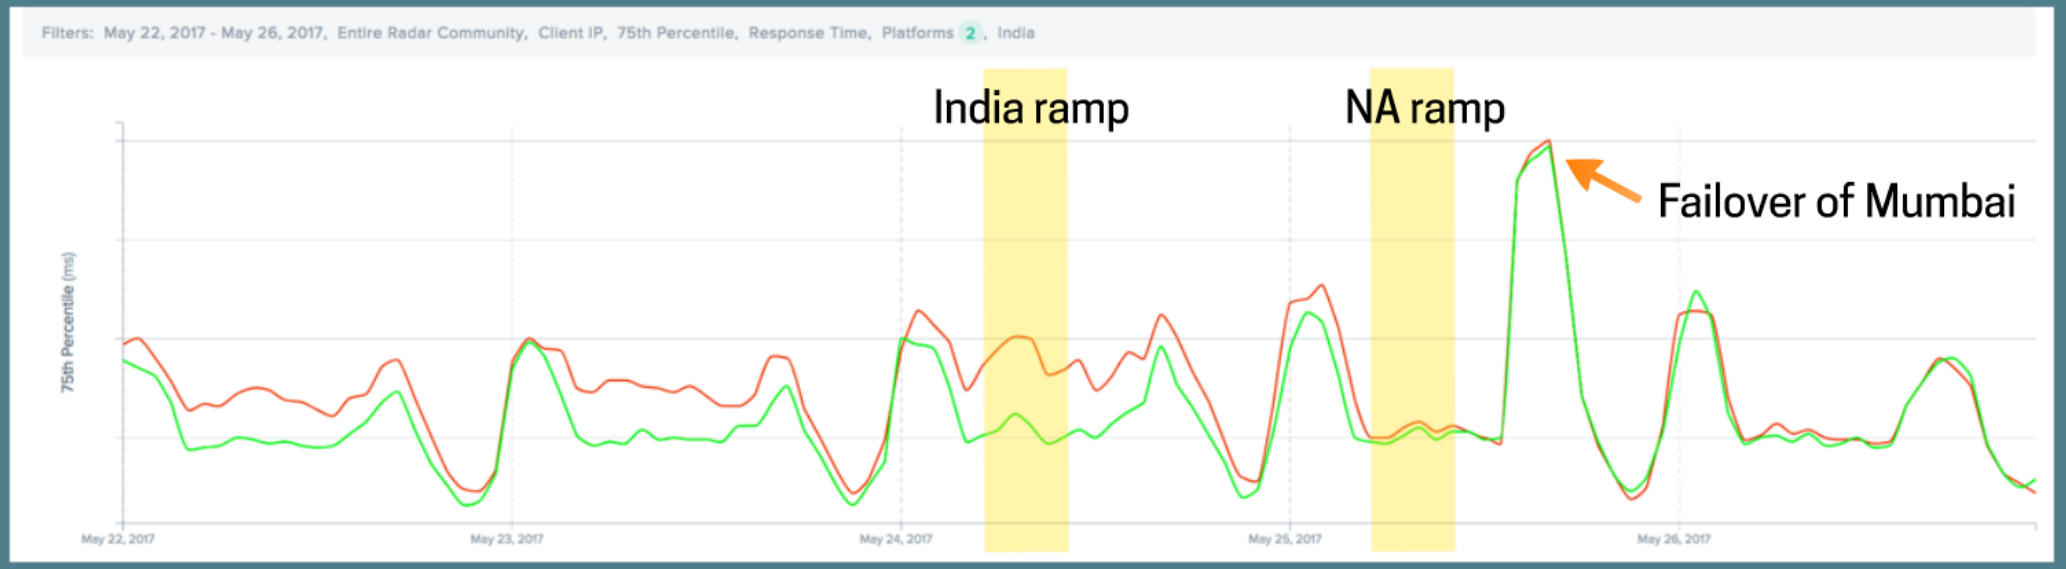

GeoDNS RUM DNS

Tata (AS4755) improves when North America ramp

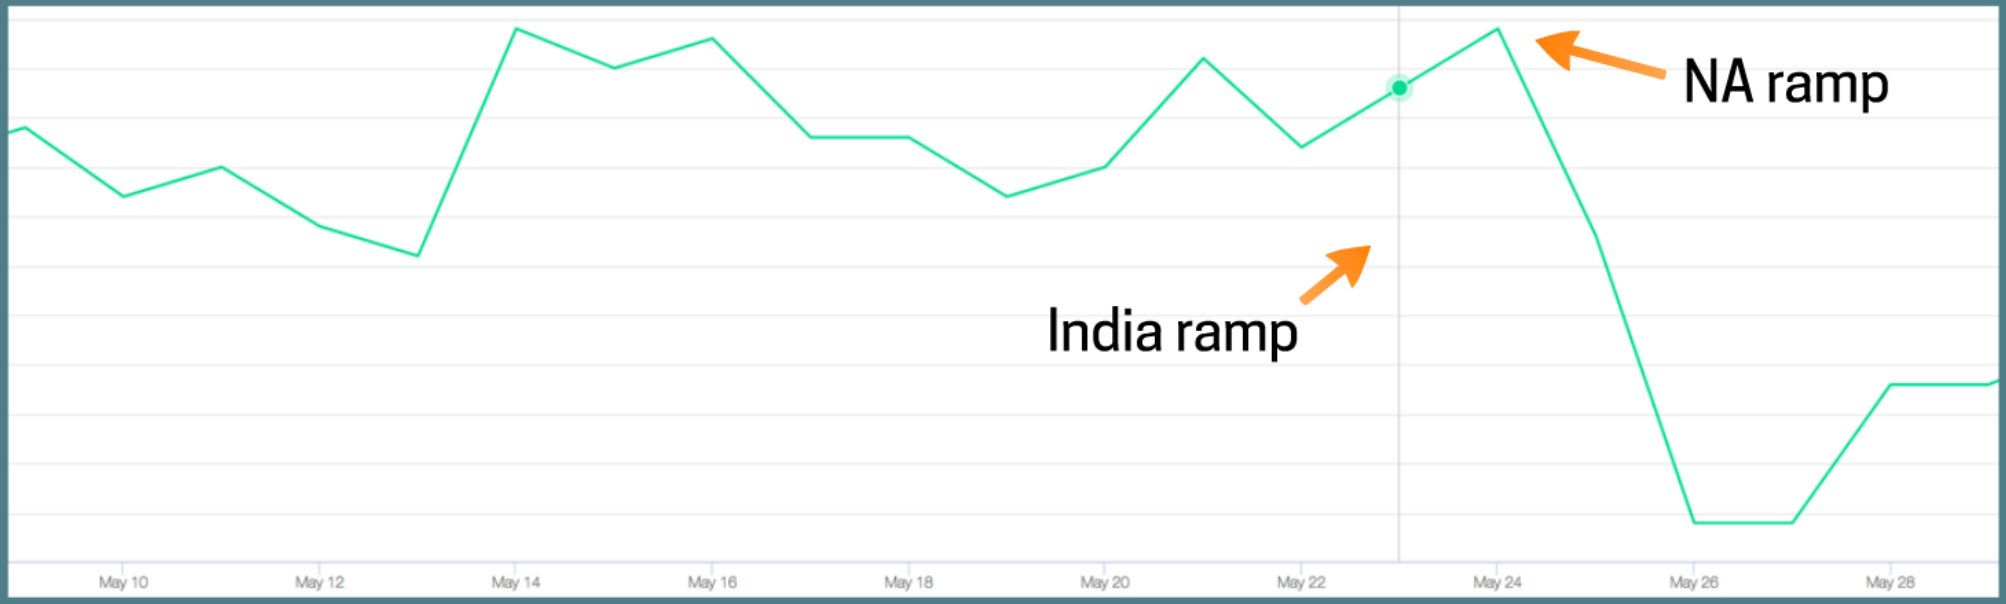

50th Connect time

Korea connect times

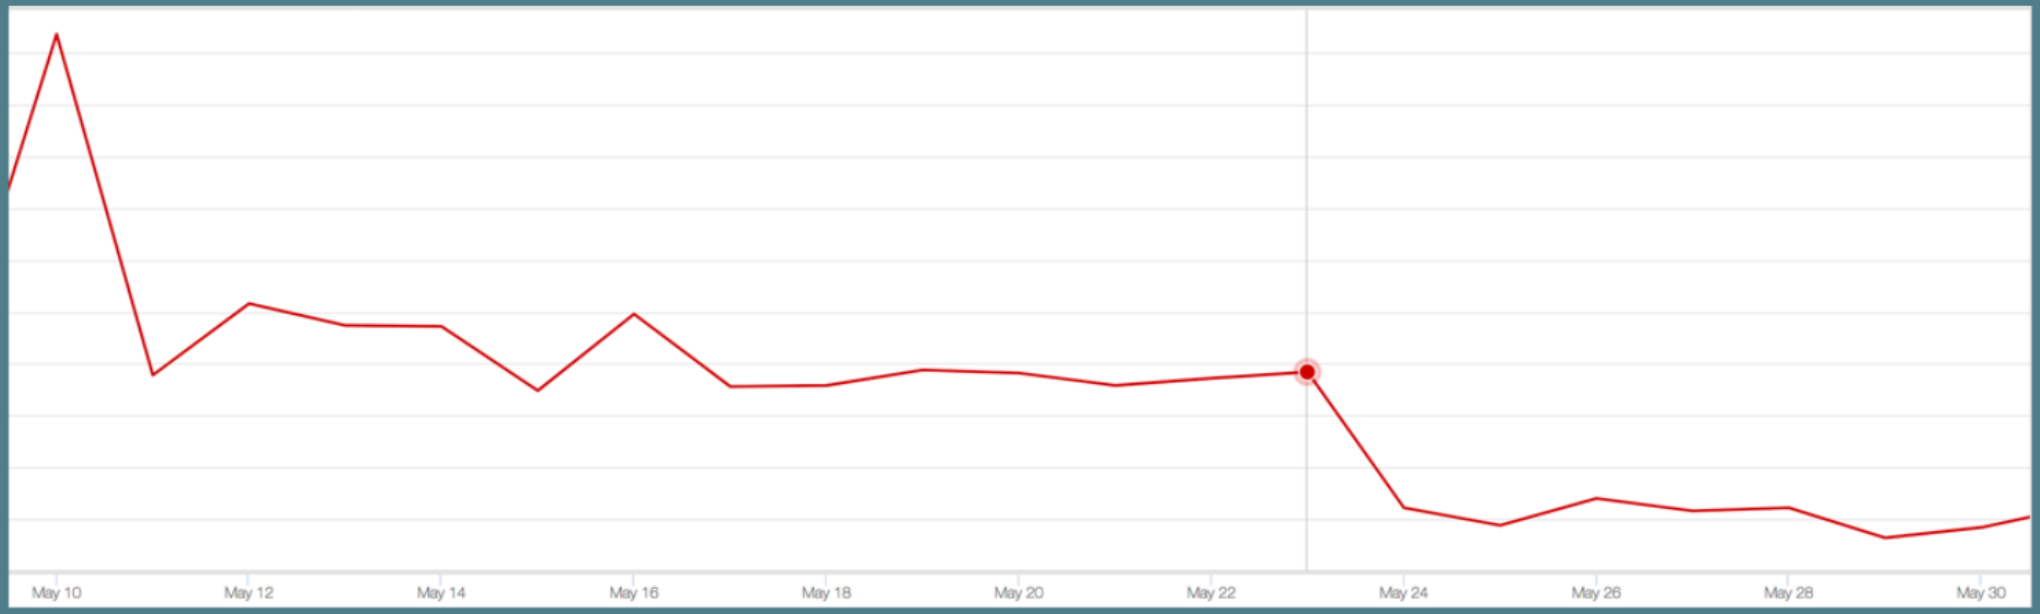

90th Connect time

RUM DNS PoP distribution for Korea

RUM DNS PoP distribution for Korea

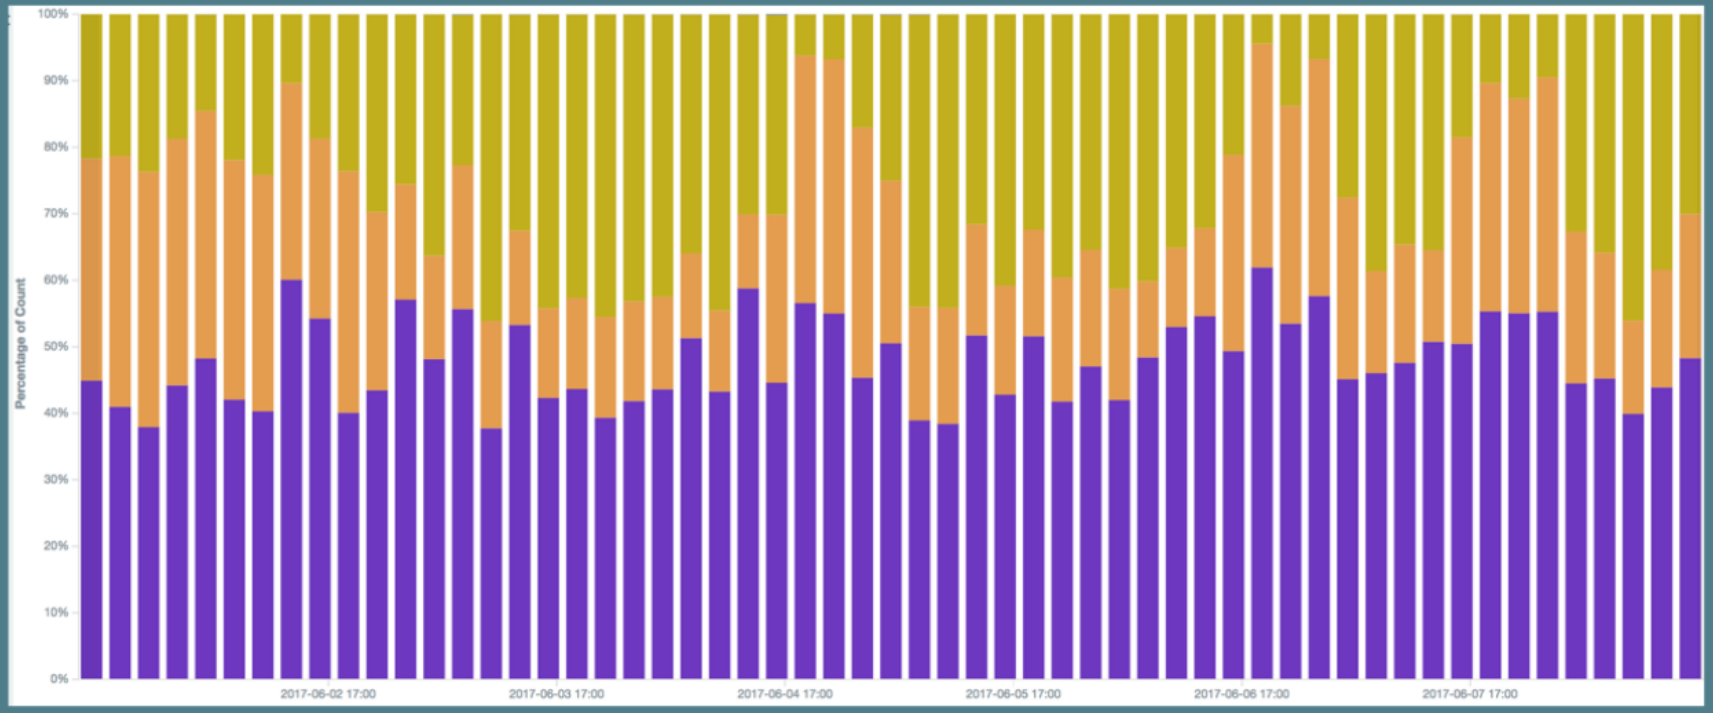

Singapore North America Hong Kong

UAE connect times

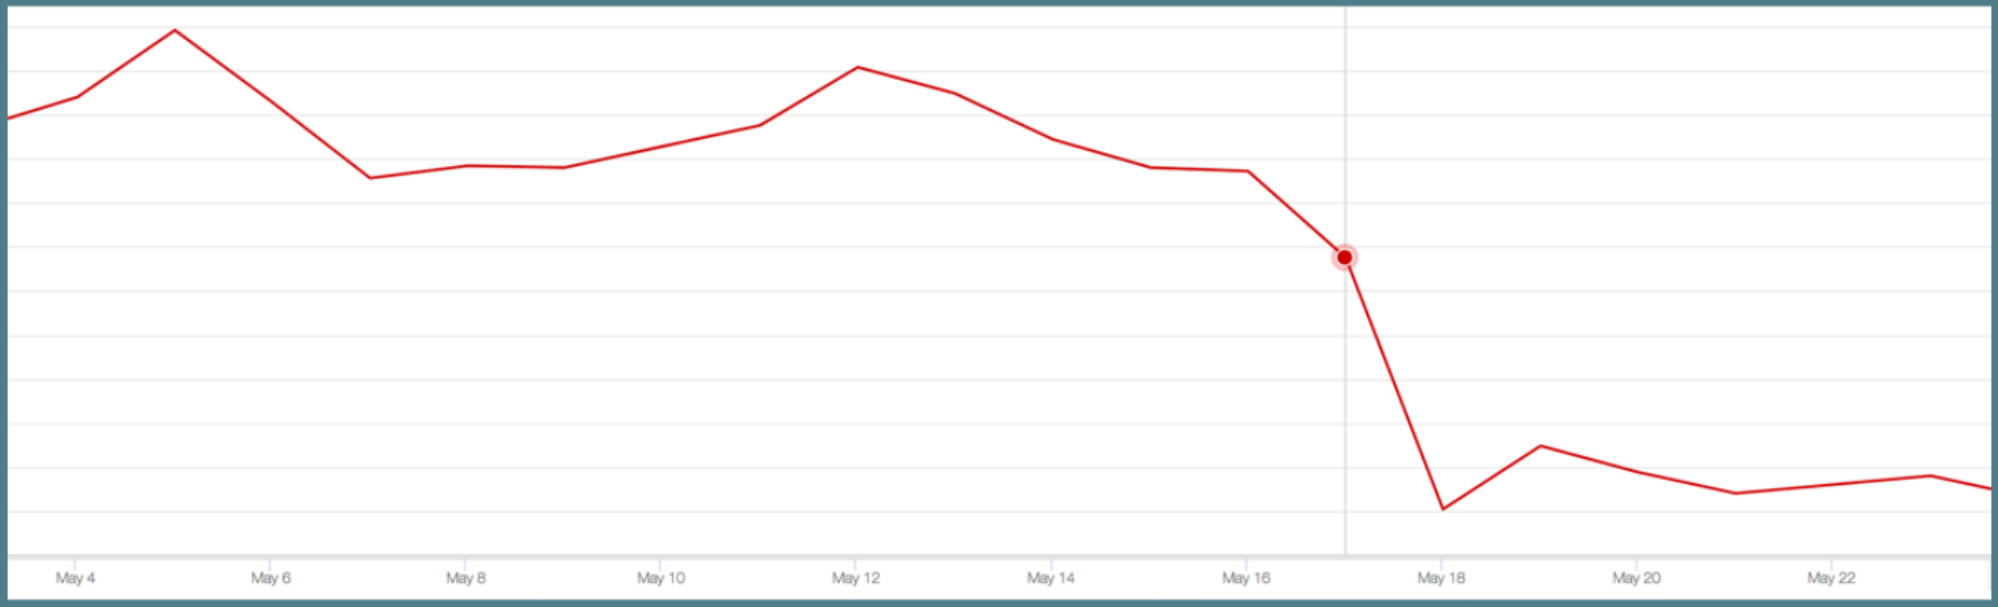

90th Connect time

DNS PoP distribution for UAE

DNS PoP distribution for UAE

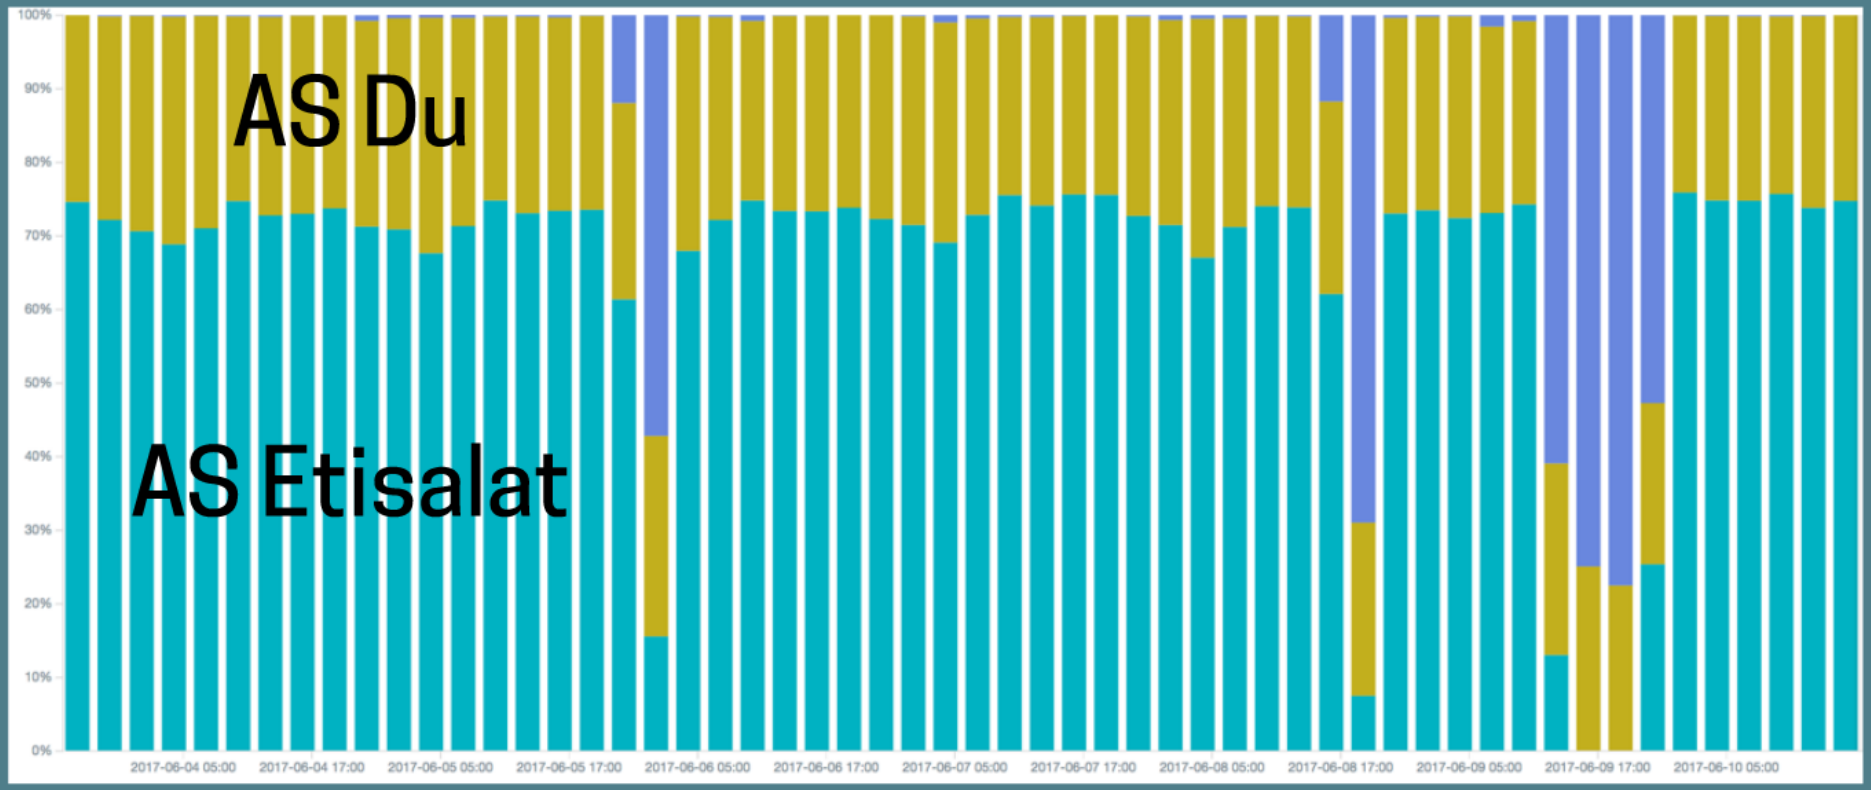

Singapore Mumbai London

RUM availabilty to Sydney PoP from various Oceania ASes

Vodafone

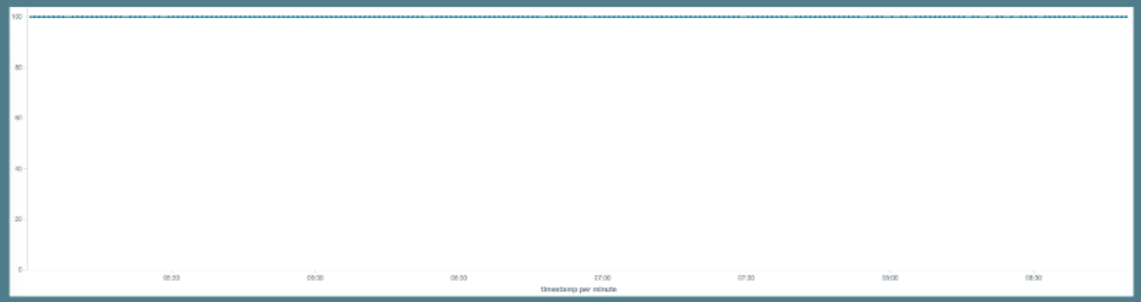

Telstra

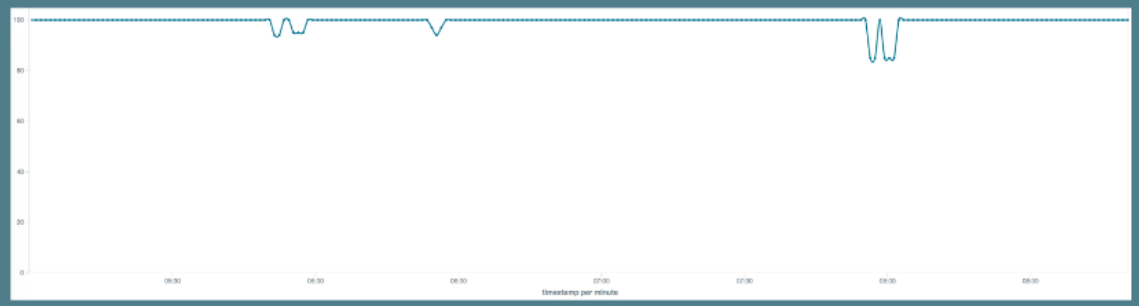

AAPT

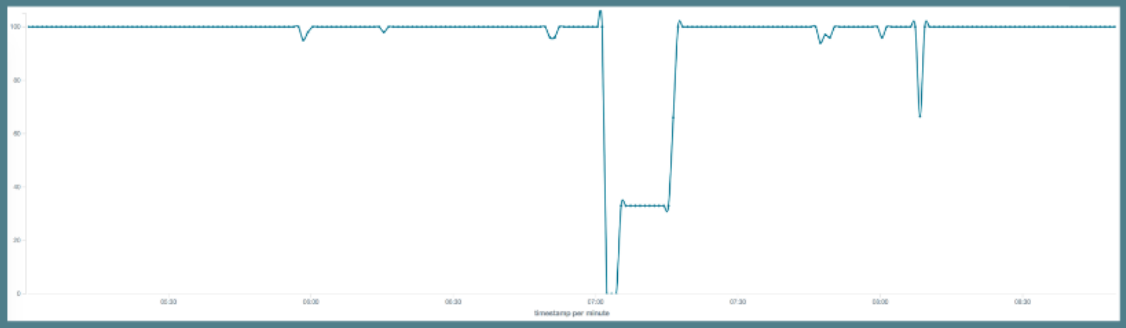

RUM availabilty for our PoPs from ASN 2764 (AAPT)

Sydney

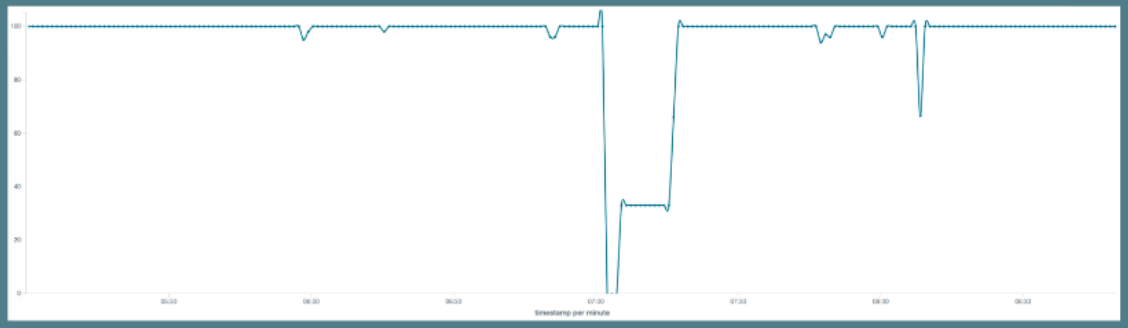

Singapore

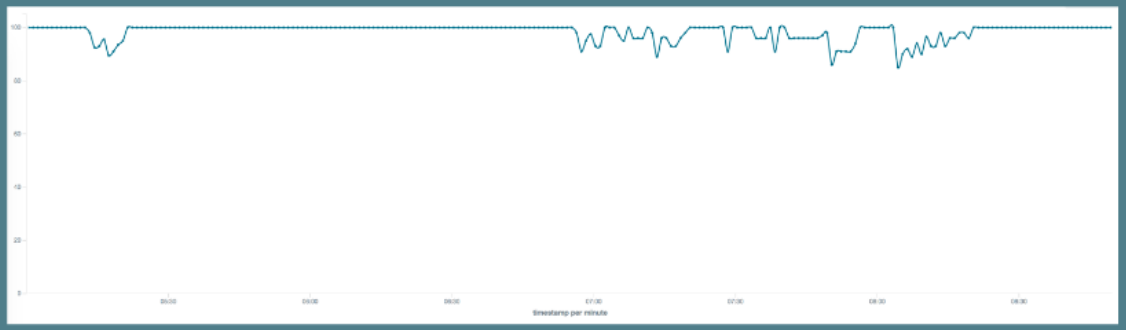

Hong Kong

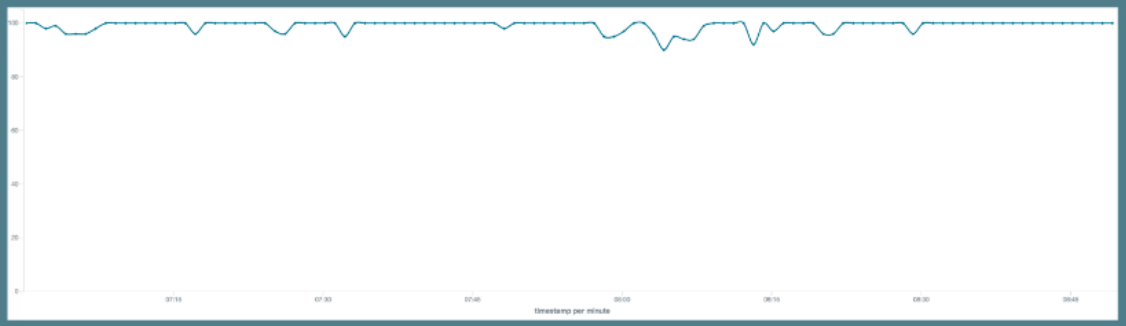

RUM availability for Sydney from ASN 2764 (AAPT)

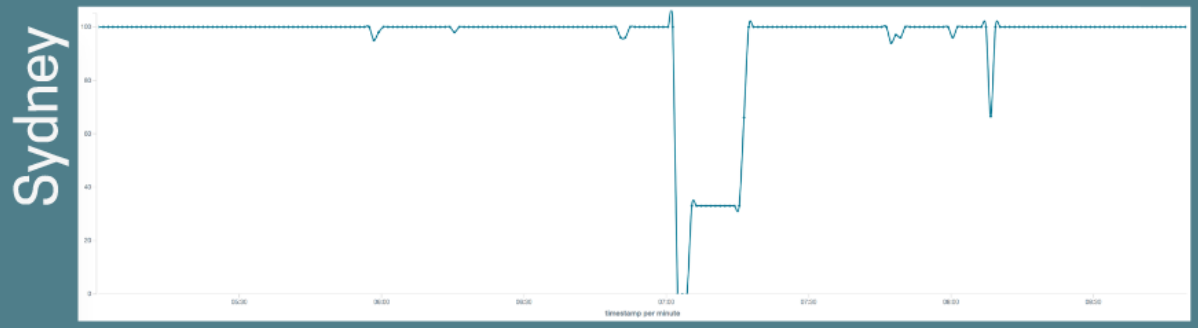

RUM DNS decisions for AAPT route around outage

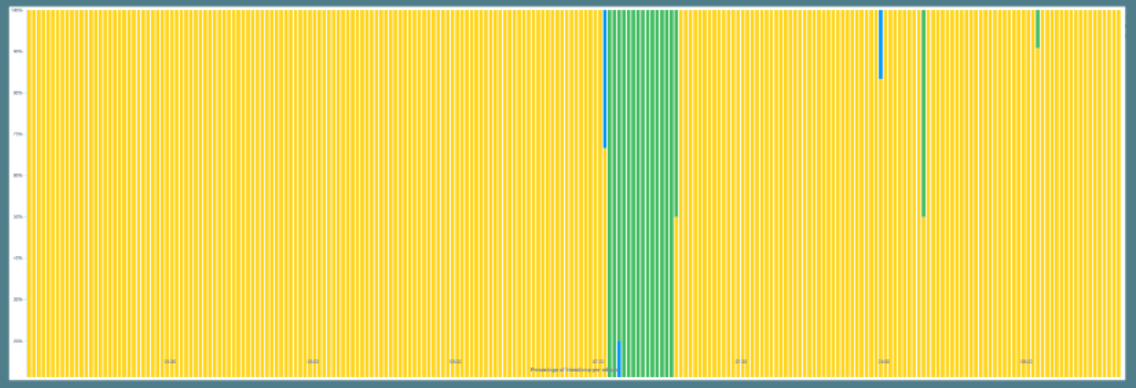

Sydney Singapore Hong Kong

RUM DNS for PoPs v2 is a blend of performance and programmability

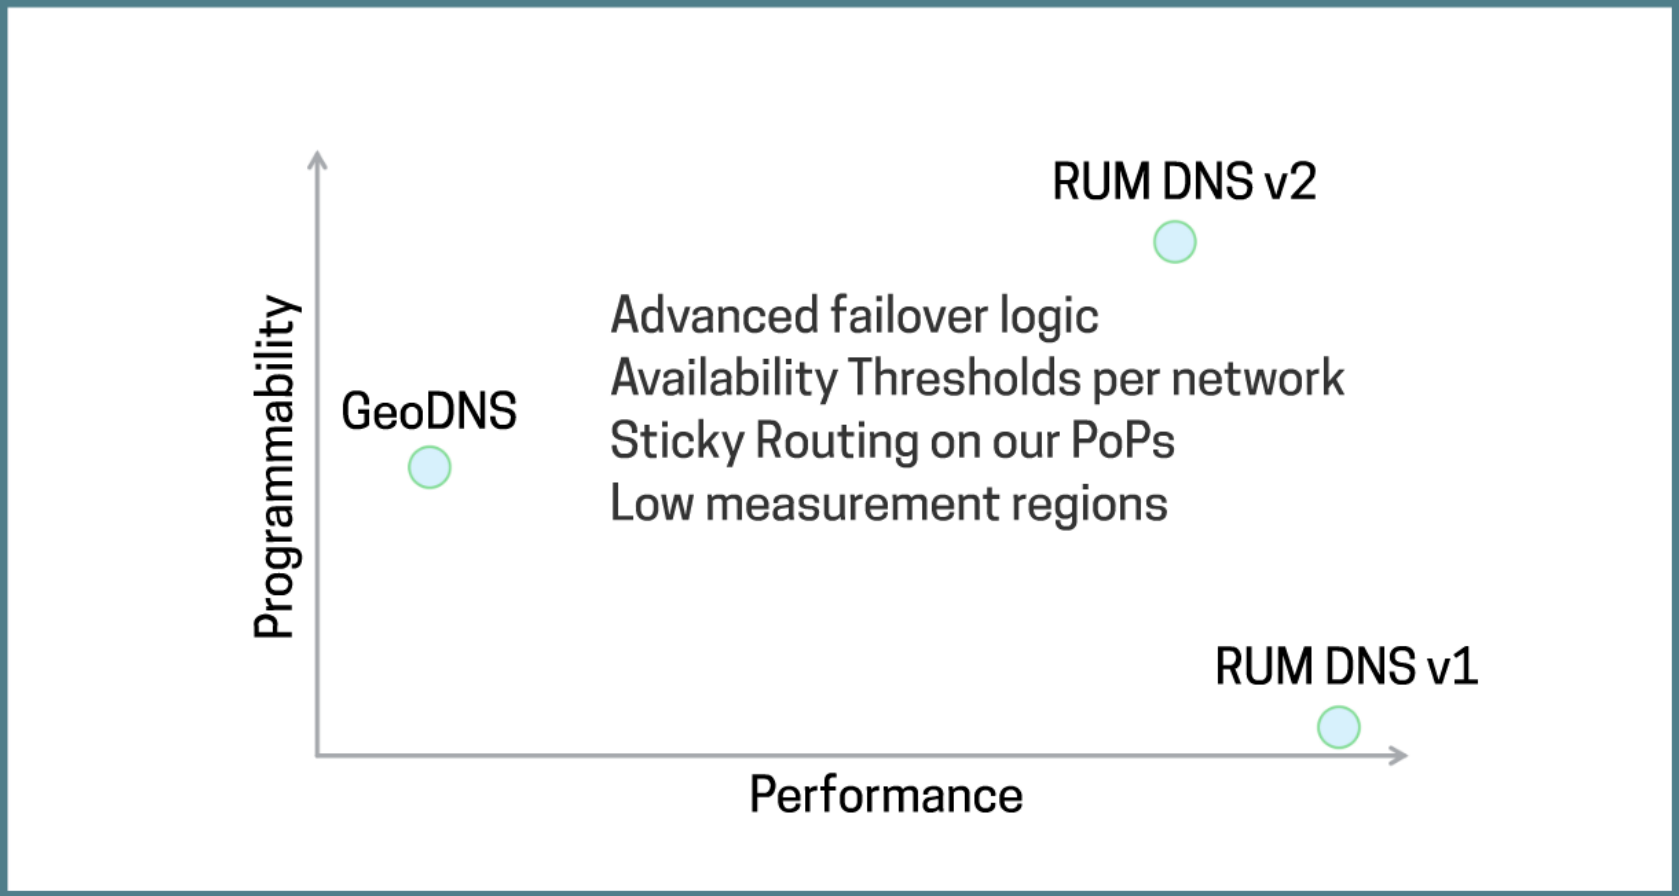

Overrides put bumpers on low measurement regions

- Assumed Africa should be sent to Europe or North America
- v2 app knows right answer
- v2 app issues reason codes when overrides aren't correct
- Every DNS answer is logged and sent to ELK

The R code: When v2 knows it's not doing the best thing

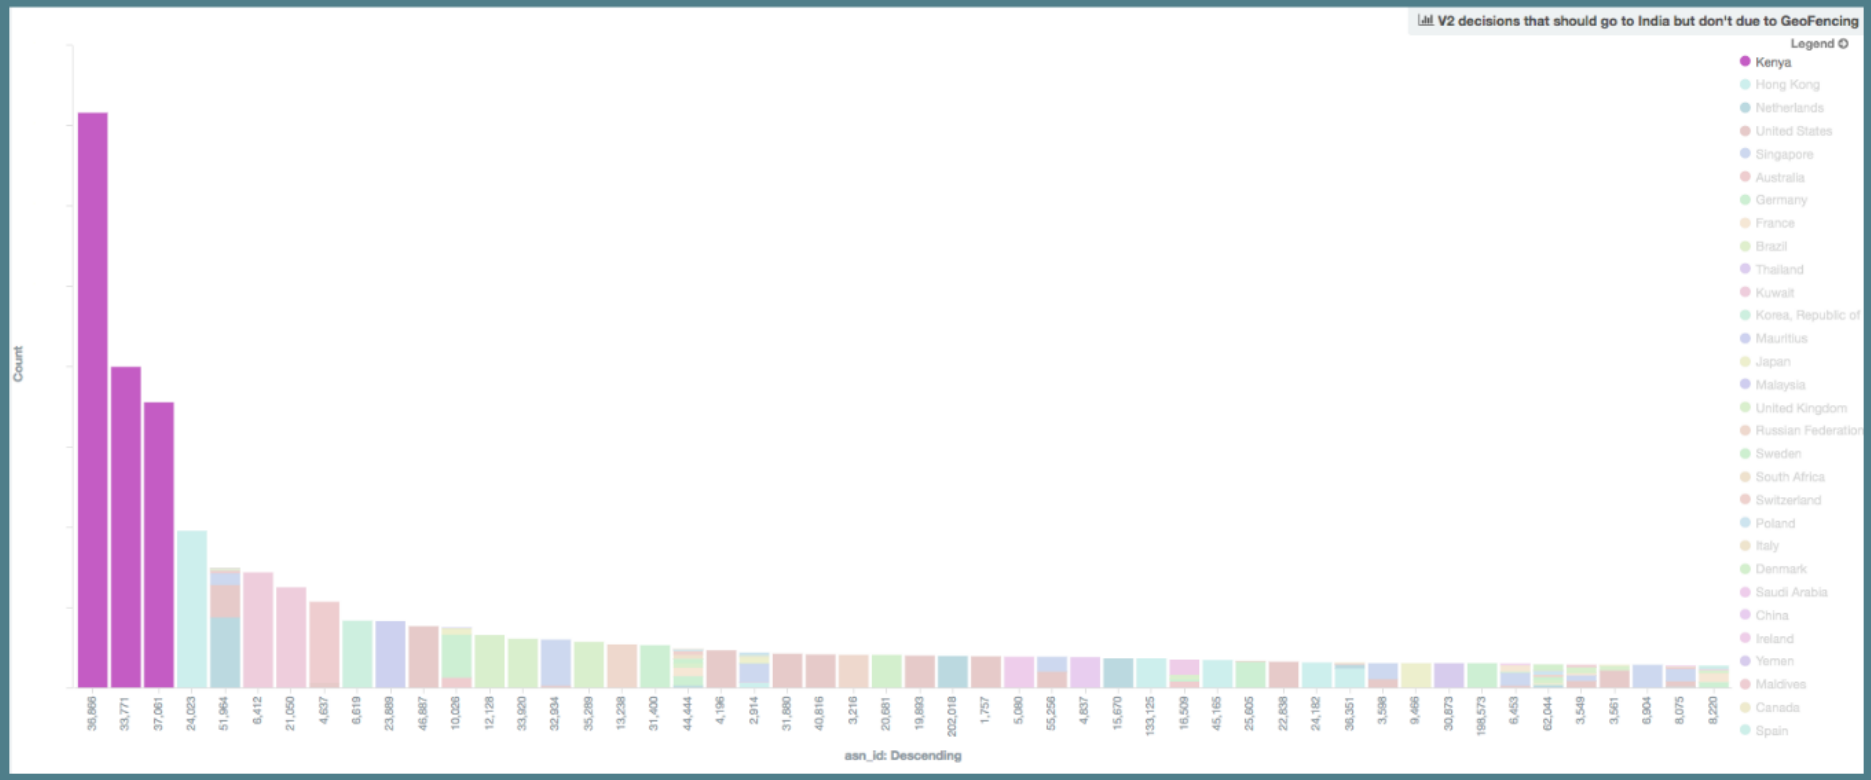

RUM DNS correctly steers Kenya to fastest PoP, in India

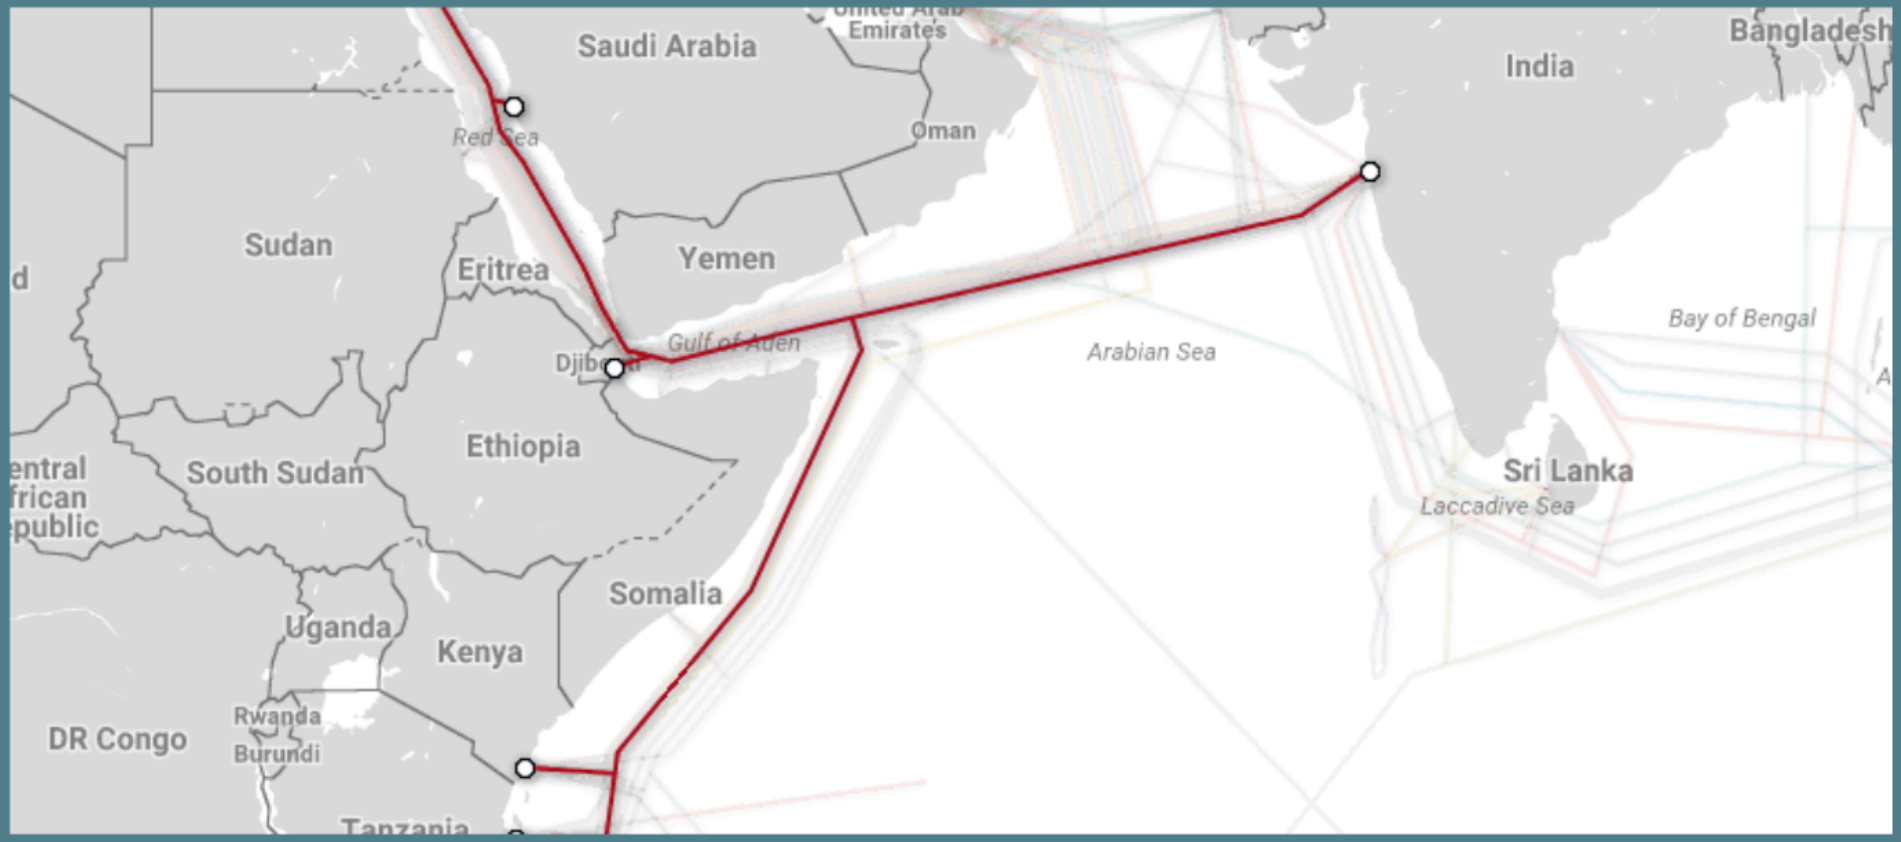

source: http://www.submarinecablemap.com/#/submarine-cable/seacomtata-tgn-eurasia

Kenyan Jamii (AS36866) better served by Mumbai

50th Connect time

Closing thoughts

Questions?