moz://a

Detecting outages with telemetry

Alessio Placitelli - @dexterp37

June 16th - Internet Measurement Village 2020

Italy, March 11th - 2020

Tales from a mid-pandemic network outage

“...failure on a foreign network...”

Source: “Sharing data on Italy’s mid-pandemic internet outage” - https://mzl.la/italy-outage

Network outage in Italy

How many Firefox desktop users were affected by the mid-pandemic outage?

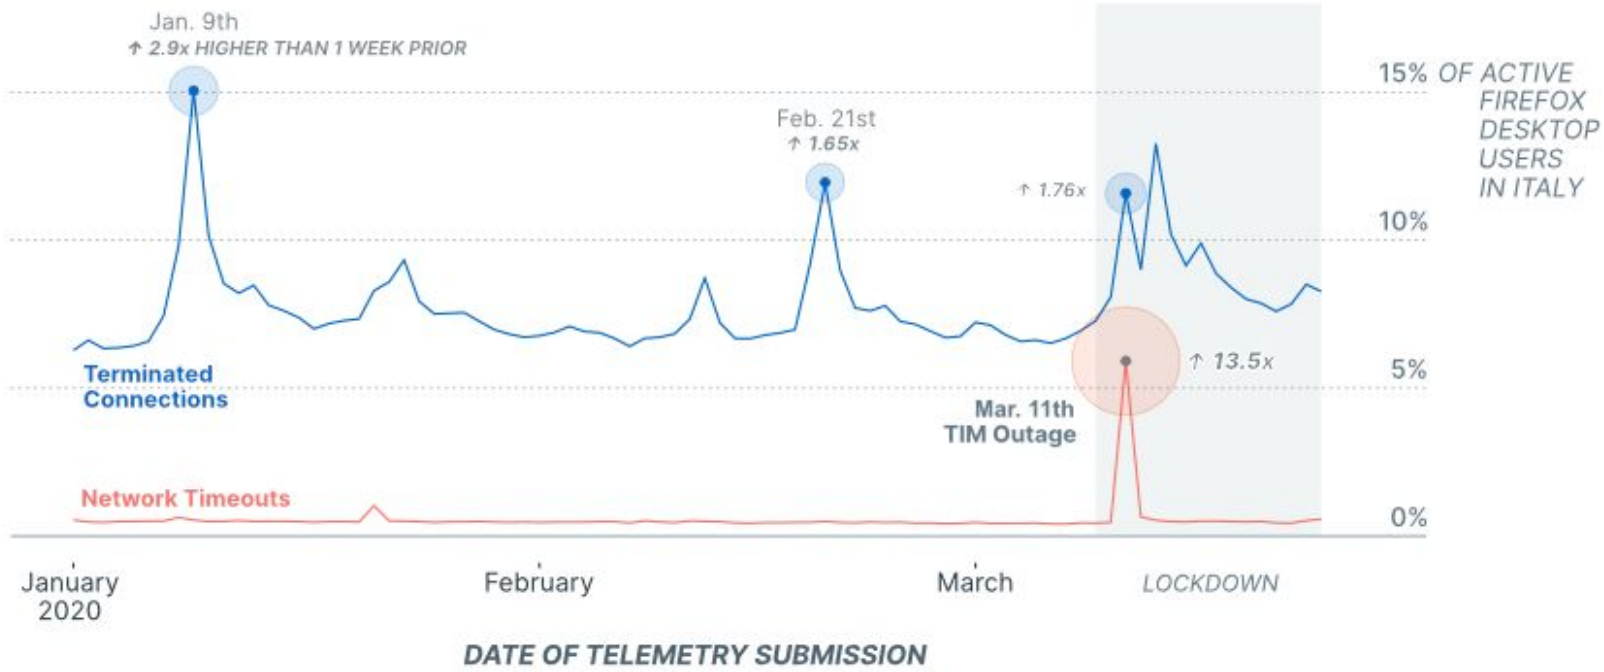

NOPE.

These were for something completely different!

“The internet is a global public resource that must remain open and accessible.”

Mozilla Manifesto

Principle 2 - https://www.mozilla.org/about/manifesto/

Key takeaways

1. Our methodology is open
2. What happened in Italy on March 11th, 2020?
3. What showed up in Jammu & Kashmir in 2019?

Telemetry

A quick overview

1. Performance **metrics** for our products
2. Packaged in **pings** sent at controlled schedules
3. Following our **Lean Data Practices** (www.leandatapractices.com)

Firefox telemetry

How does it work?

1. Relevant metrics travel in the **main** and **health** pings.

2. Documentation for metrics and pings is publicly available.

3. **probes.telemetry.mozilla.org**

The “main” ping

Schedule and properties

1. Ideally sent **once per day** around local midnight.
2. Is the main transport for Firefox telemetry.
3. Includes DNS, SSL and TLS metrics…

The “main” ping

Interesting metrics

1. dns_failed_lookup_time
2. dns_lookup_time
3. ssl_cert_verification_errors
4. http_page_tls_handshake
5. ...

The “health” ping

Schedule and properties

1. Telemetry health about… telemetry.
2. Extremely small (~800 bytes).
3. Collected **at most** once per hour in case of problems.
4. Includes the reason why the HTTPS upload failed.

Our open methodology

From raw data to pretty graphs

Throw away that IP address!

Right after matching the IP with a country lookup, at ingestion!

https://github.com/mozilla/gcp-ingestion/blob/fbfb5d28490a17d43329b44a1a8259bbcc0d7b20/ingestion-beam/src/main/java/com/mozilla/telemetry/Decoder.java#L64-L69

Cleanup:
remove “inactive” sessions

Not all the “main” pings are representative.

“Who can even open 100 websites in 1 second?”

Aggregation:
step 1 - geographical

Group the data by Country.

Drop the data for Countries with too few samples.

Aggregation:
step 2 - counting things!

Count how many sessions reported a metric, within the given timeframe.

Example: how many sessions had **DNS_LOOKUP_TIME**?

Aggregation:
step 3 - create timing profiles

Combine the user-reported time distributions in a single distribution, for a given timeframe.

Example: what's the shape of **DNS_LOOKUP_TIME** in Italy, today?

Investigation:
look for anomalies in the data

How do certain measures compare against a baseline?

Were there anomalous spikes, surges, holes in the time series?

Jammu & Kashmir - 2019

Network interferences starting from August 5th

Jammu & Kashmir

How many Firefox desktop users were affected (normalized count)?

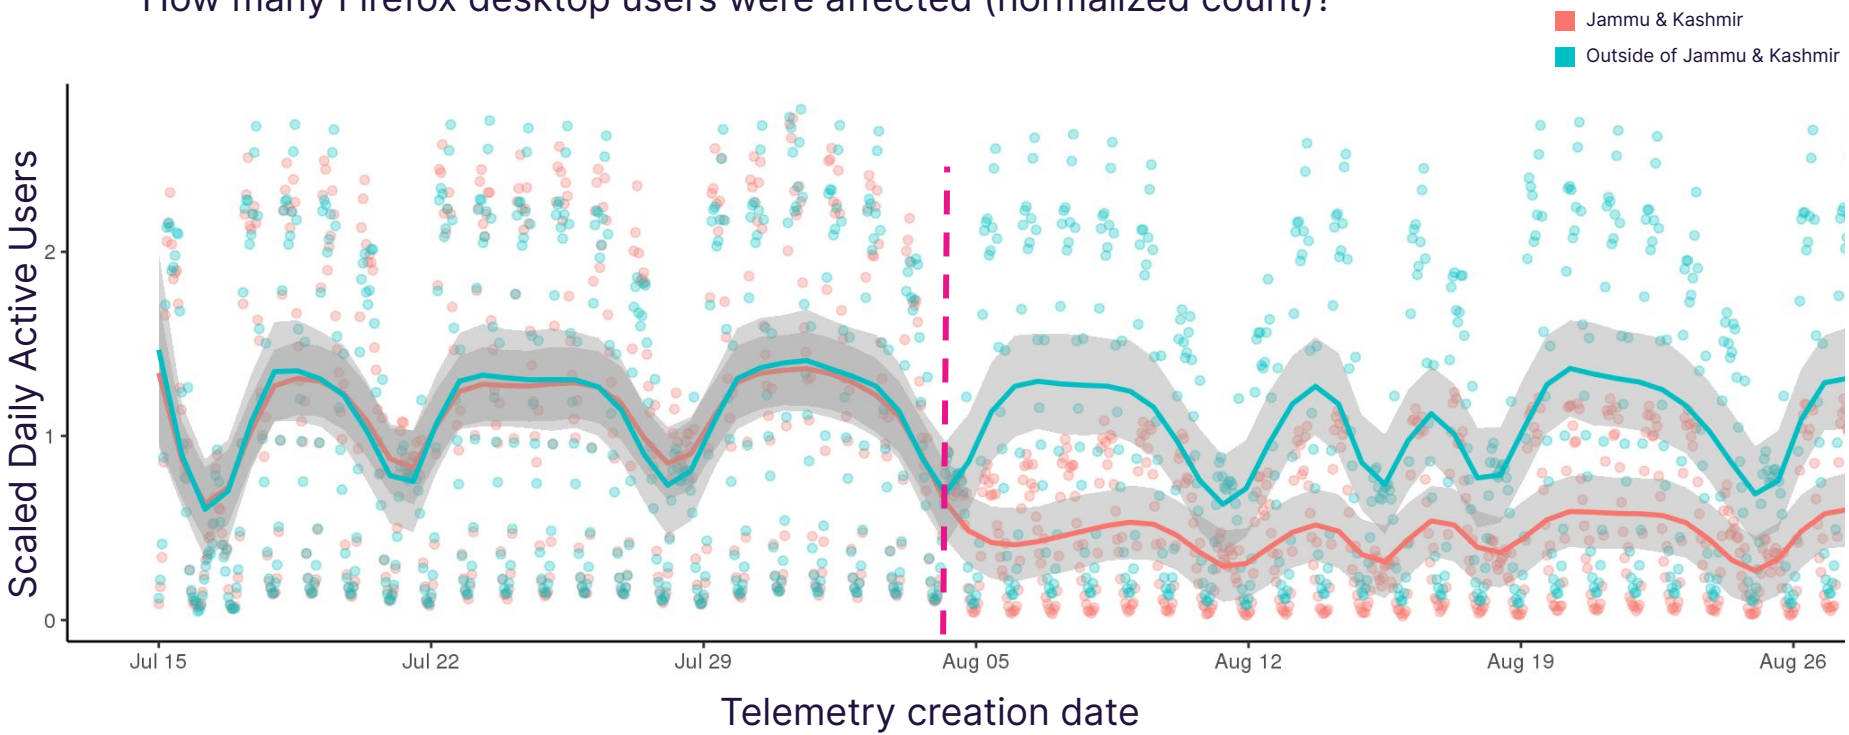

Jammu & Kashmir

The average time it takes for an ***unsuccessful*** DNS resolution, in milliseconds

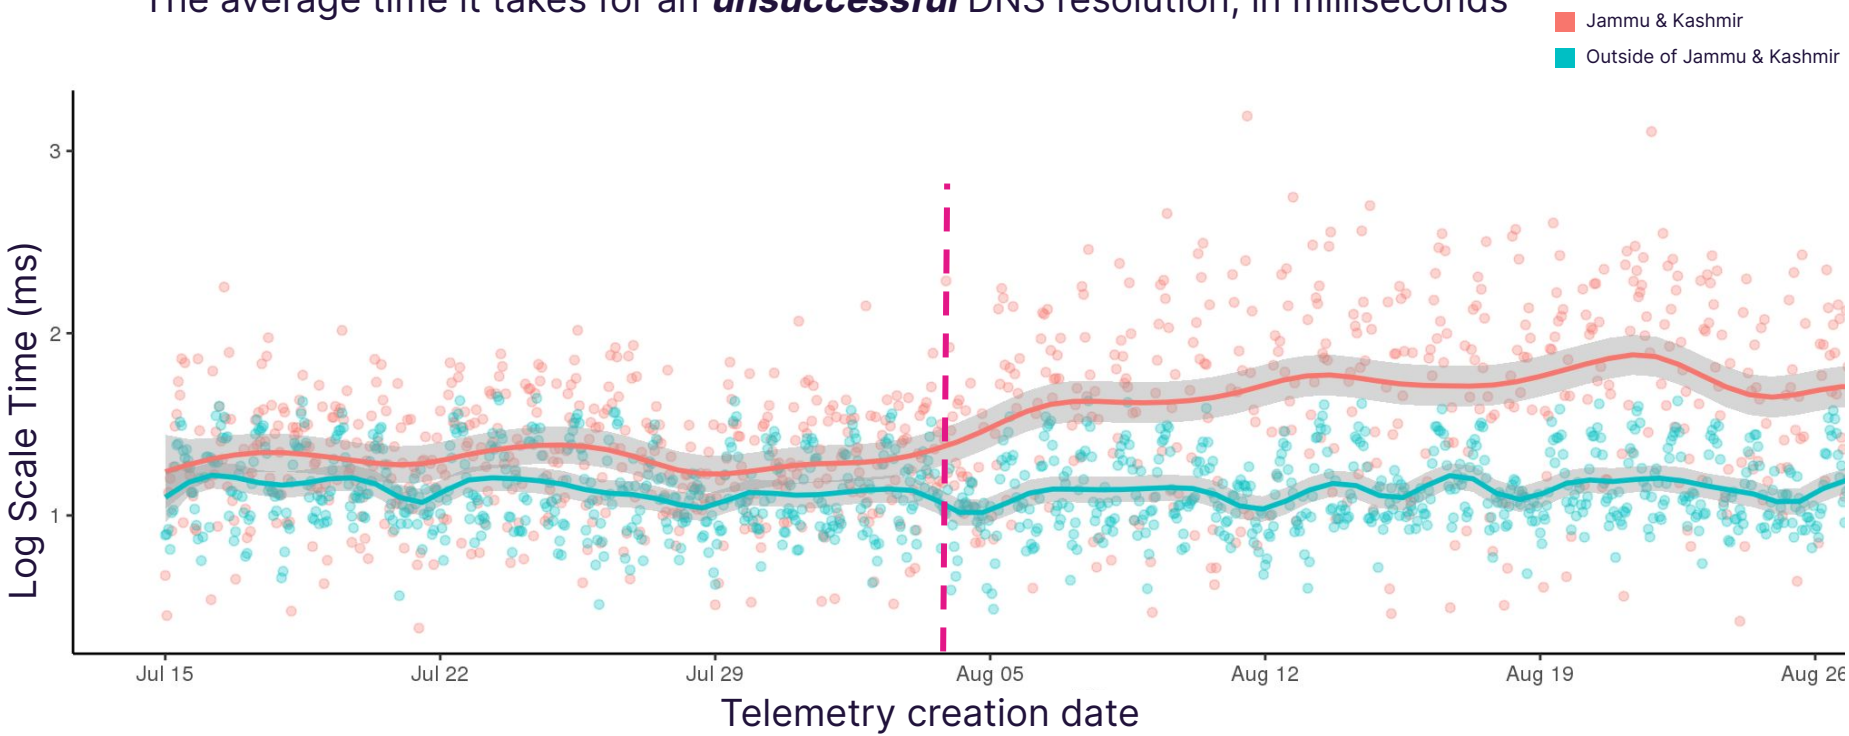

Jammu & Kashmir

The proportion of active session with ***no DNS resolved***

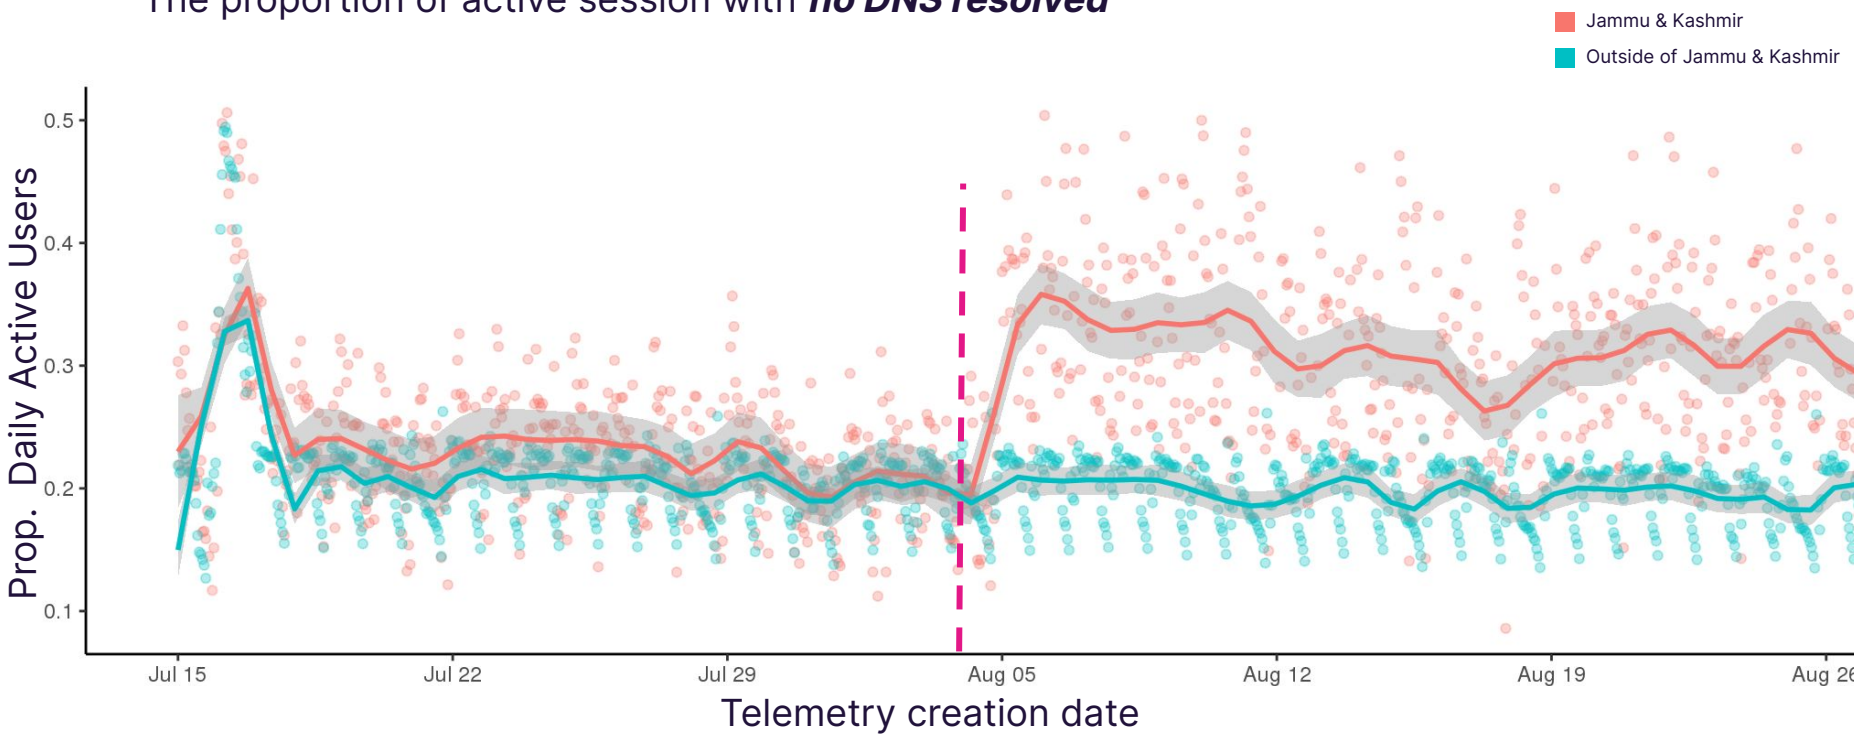

What’s next?

How are we moving this project forward

Productionize our datasets

01

Validate the data

02

Community collaboration

03

Our team

Solana Larsen
Editor, Internet Health Report

Saptarshi Guha
Data Scientist

Jochai Ben-Avie
Head of International Public Policy

Alessio Placitelli
Telemetry Engineer, Project Lead

Special thanks to **Rebecca Weiss** for advising on the project, and to **Hamilton Ulmer** for the graphics on the Italian focus

Thank you!

Reach out to: **outages@mozilla.com**