



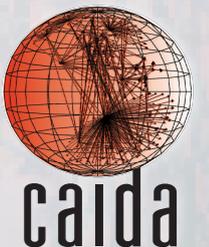
MEASUREMENT AND ANALYSIS OF INTERNET CONGESTION

Amogh Dhamdhere
amogh@caida.org

with D. Clark, k Claffy, M. Luckie, B. Huffaker, R. Mok,
A. Gamero-Garrido, A. Snoeren, G. Akiwate, V. Bajpai, R.
Fanou, E. Carisimo



MIT



Interdomain Congestion

- We are developing methods to measure the location and extent of interdomain congestion
- **Our goals (1) a system to monitor interdomain links and their congestion state, (2) a “congestion heat map” of the Internet, (3) increasing transparency, empirical grounding of debate**
- Part of a 3 year NSF-funded project on topology and congestion measurement

In the Press



France Telecom Accused Of Holding YouTube Videos Hostage Unless It Gets More Money

from the *more-peering-disputes* dept

'Peering' Into AOL-MSN Outage

Level 3 and Comcast Issue Statement

Jul 16, 2013

BROOMFIELD, Colo., July 16, 2013 – Level 3 and Comcast have resolved their prior interconnect dispute on mutually satisfactory terms. Details will not be released.

Confirmed: Comcast and Netflix have signed a paid peering agreement

Netflix packets being dropped every day because Verizon wants more money

Verizon wants to be paid by consumers and Cogent, but Cogent refuses to pay.

Cogent Gearing for Another Peering Battle

by [Stacey Higginbotham](#) FEB. 23, 2014 - 9:27 AM

Verizon denies using net neutrality victory to sabotage Netflix, Amazon

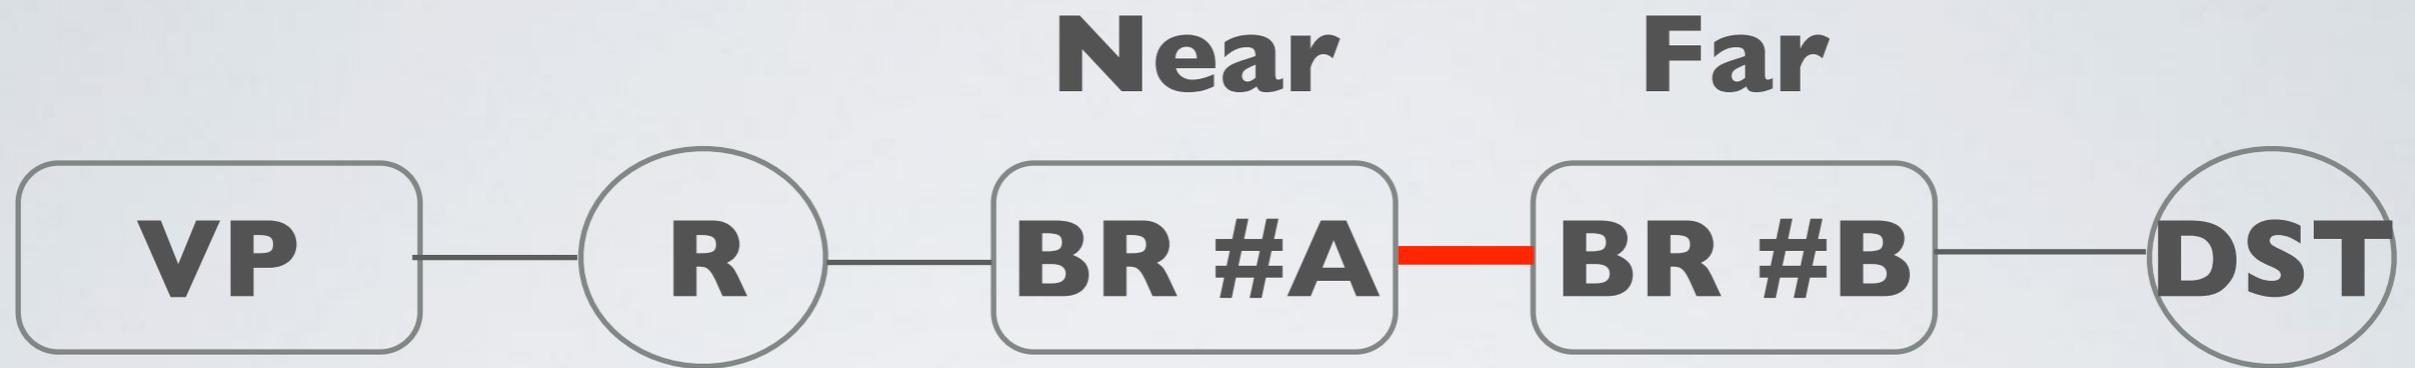
BY **BRIAN FUNG** February 5 at 1:59 pm

Netflix still sucks on AT&T, and now AT&T plans to offer Netflix clone

AT&T partners with an investment firm to buy and launch streaming services.



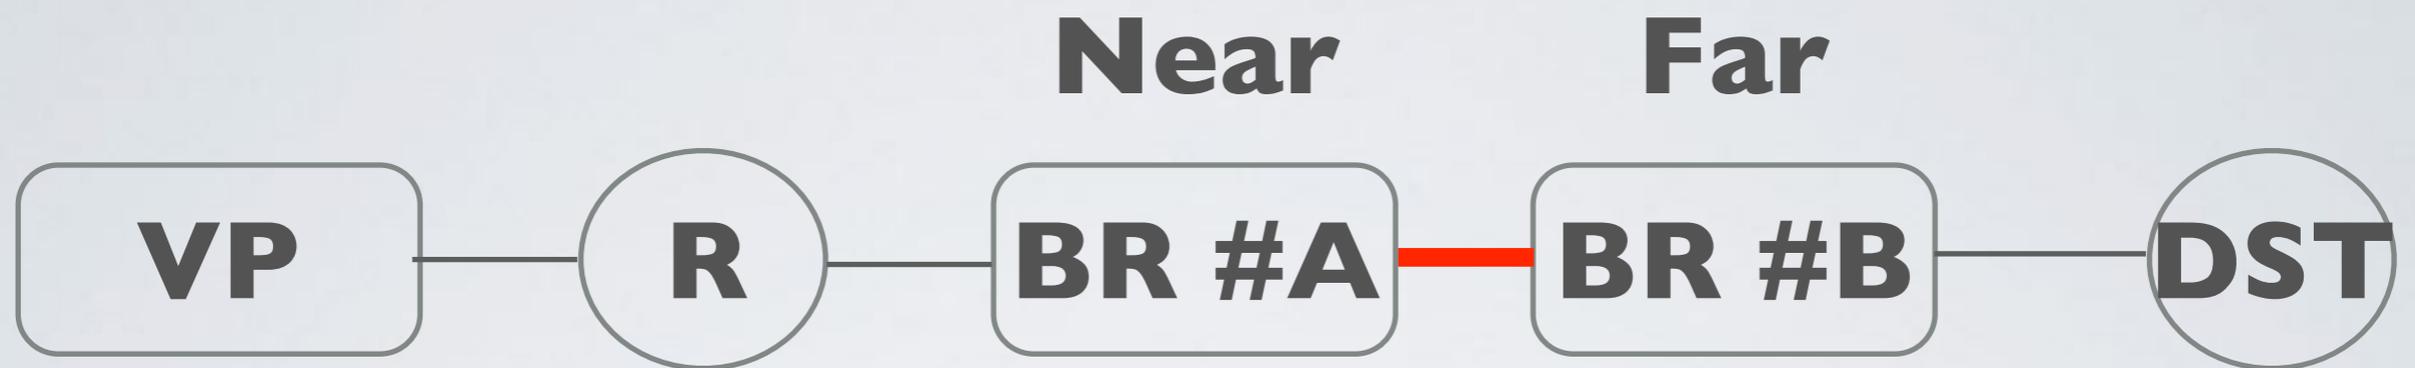
Method: Time Series Latency Probing



Vantage Point

**Border Routers on
Interesting Link**

Method: Time Series Latency Probing

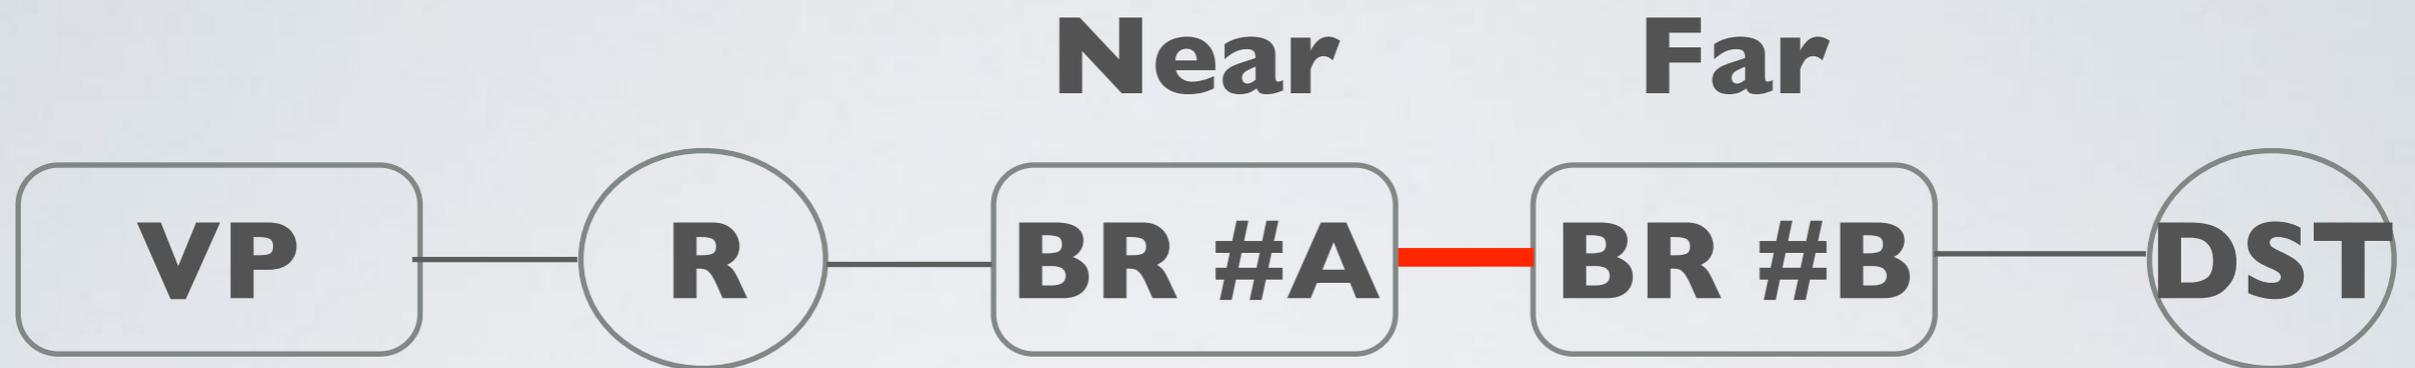


Vantage Point

**Border Routers on
Interesting Link**

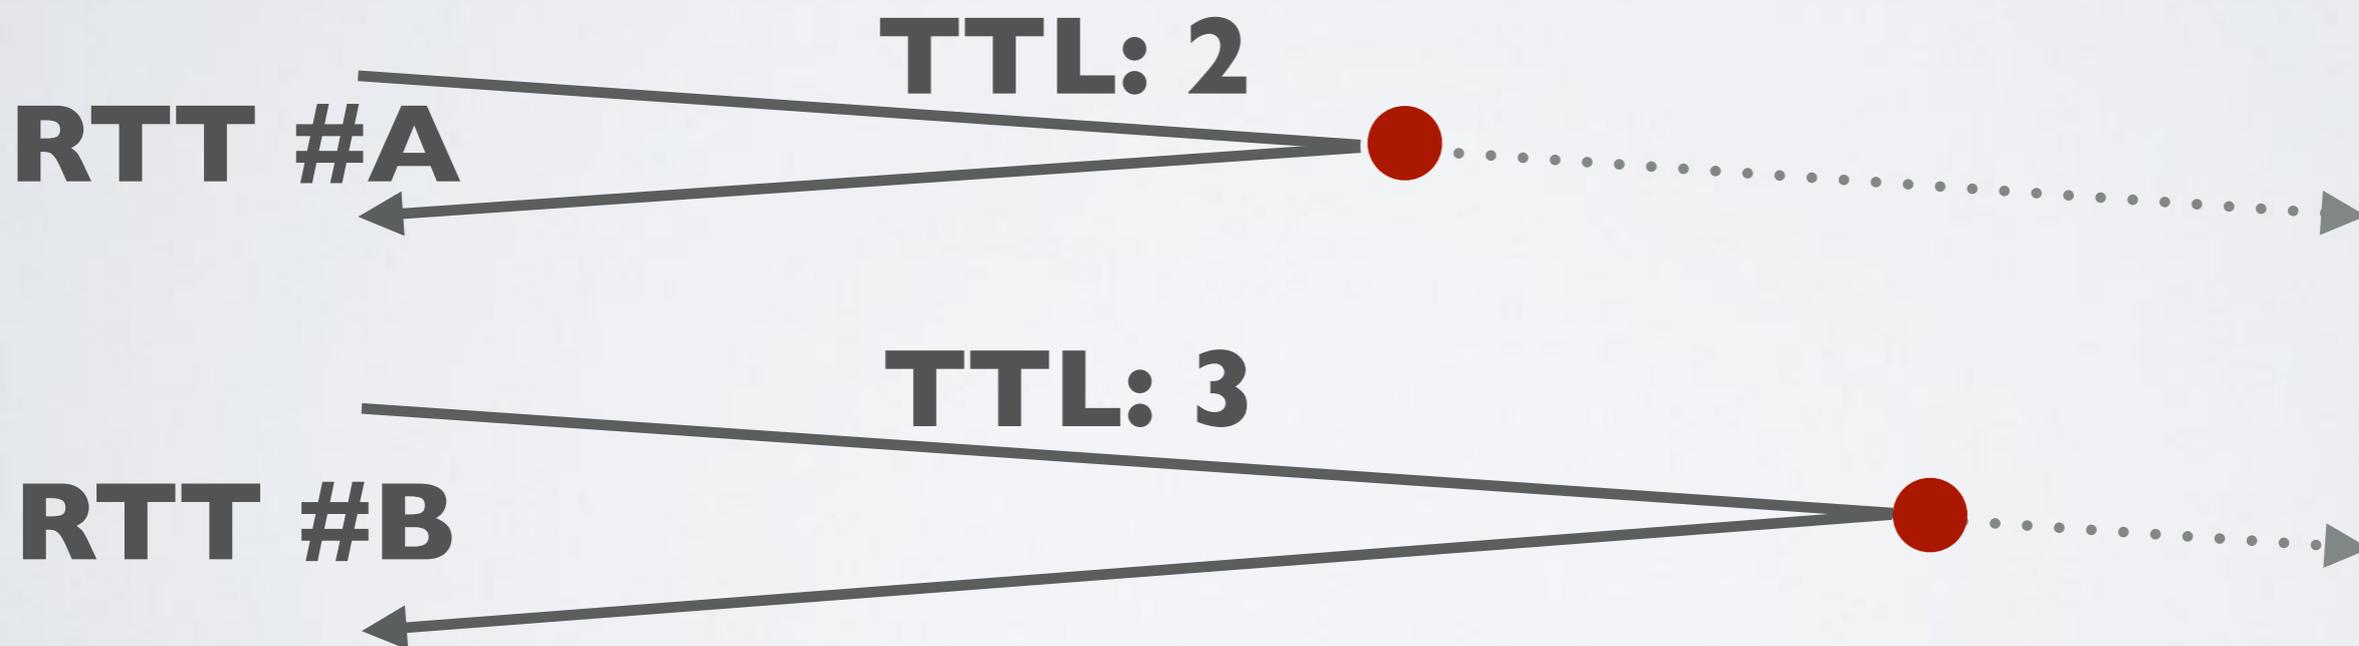


Method: Time Series Latency Probing

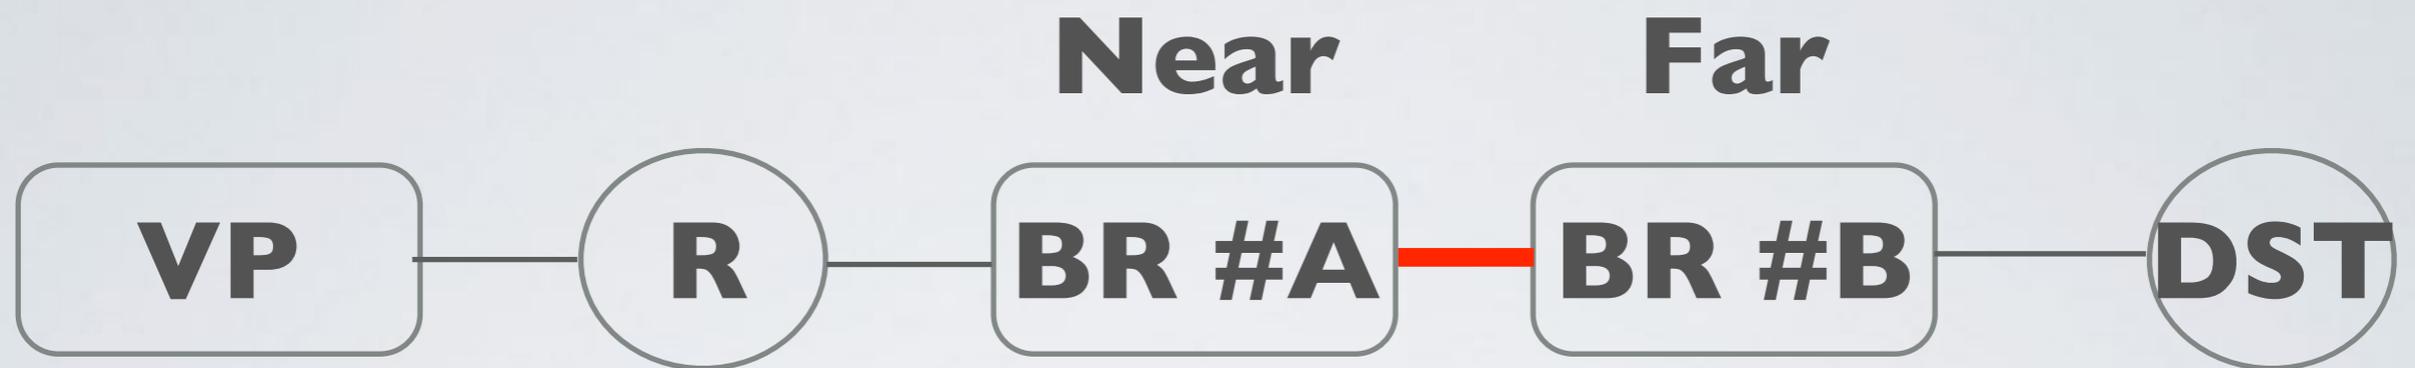


Vantage Point

**Border Routers on
Interesting Link**

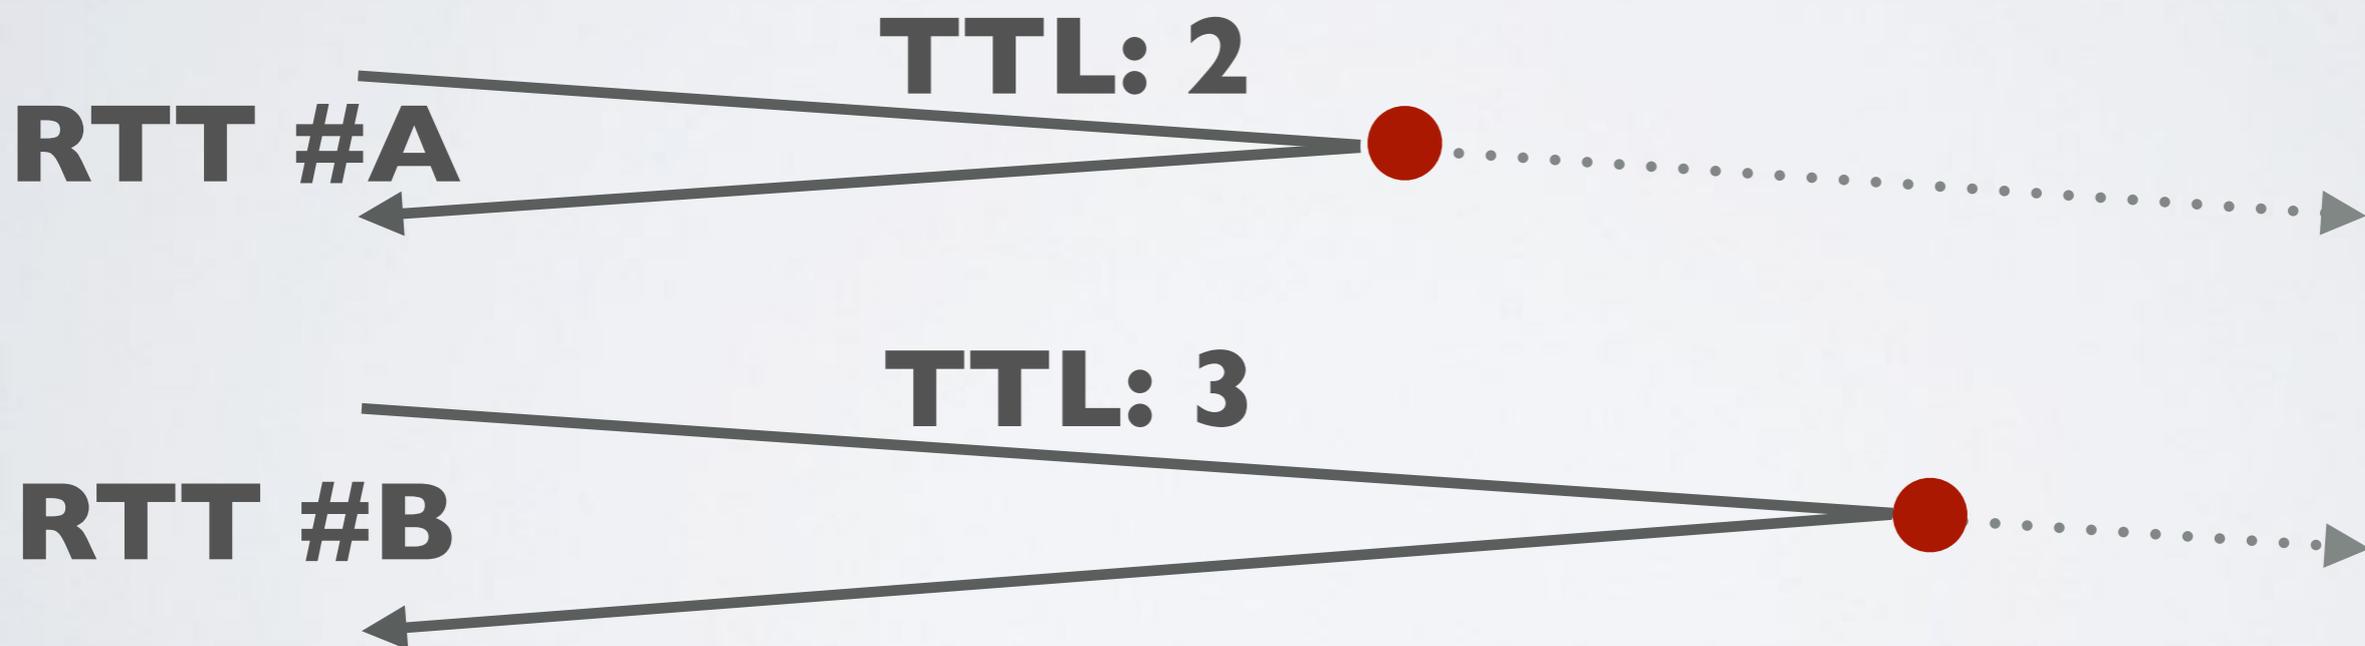


Method: Time Series Latency Probing



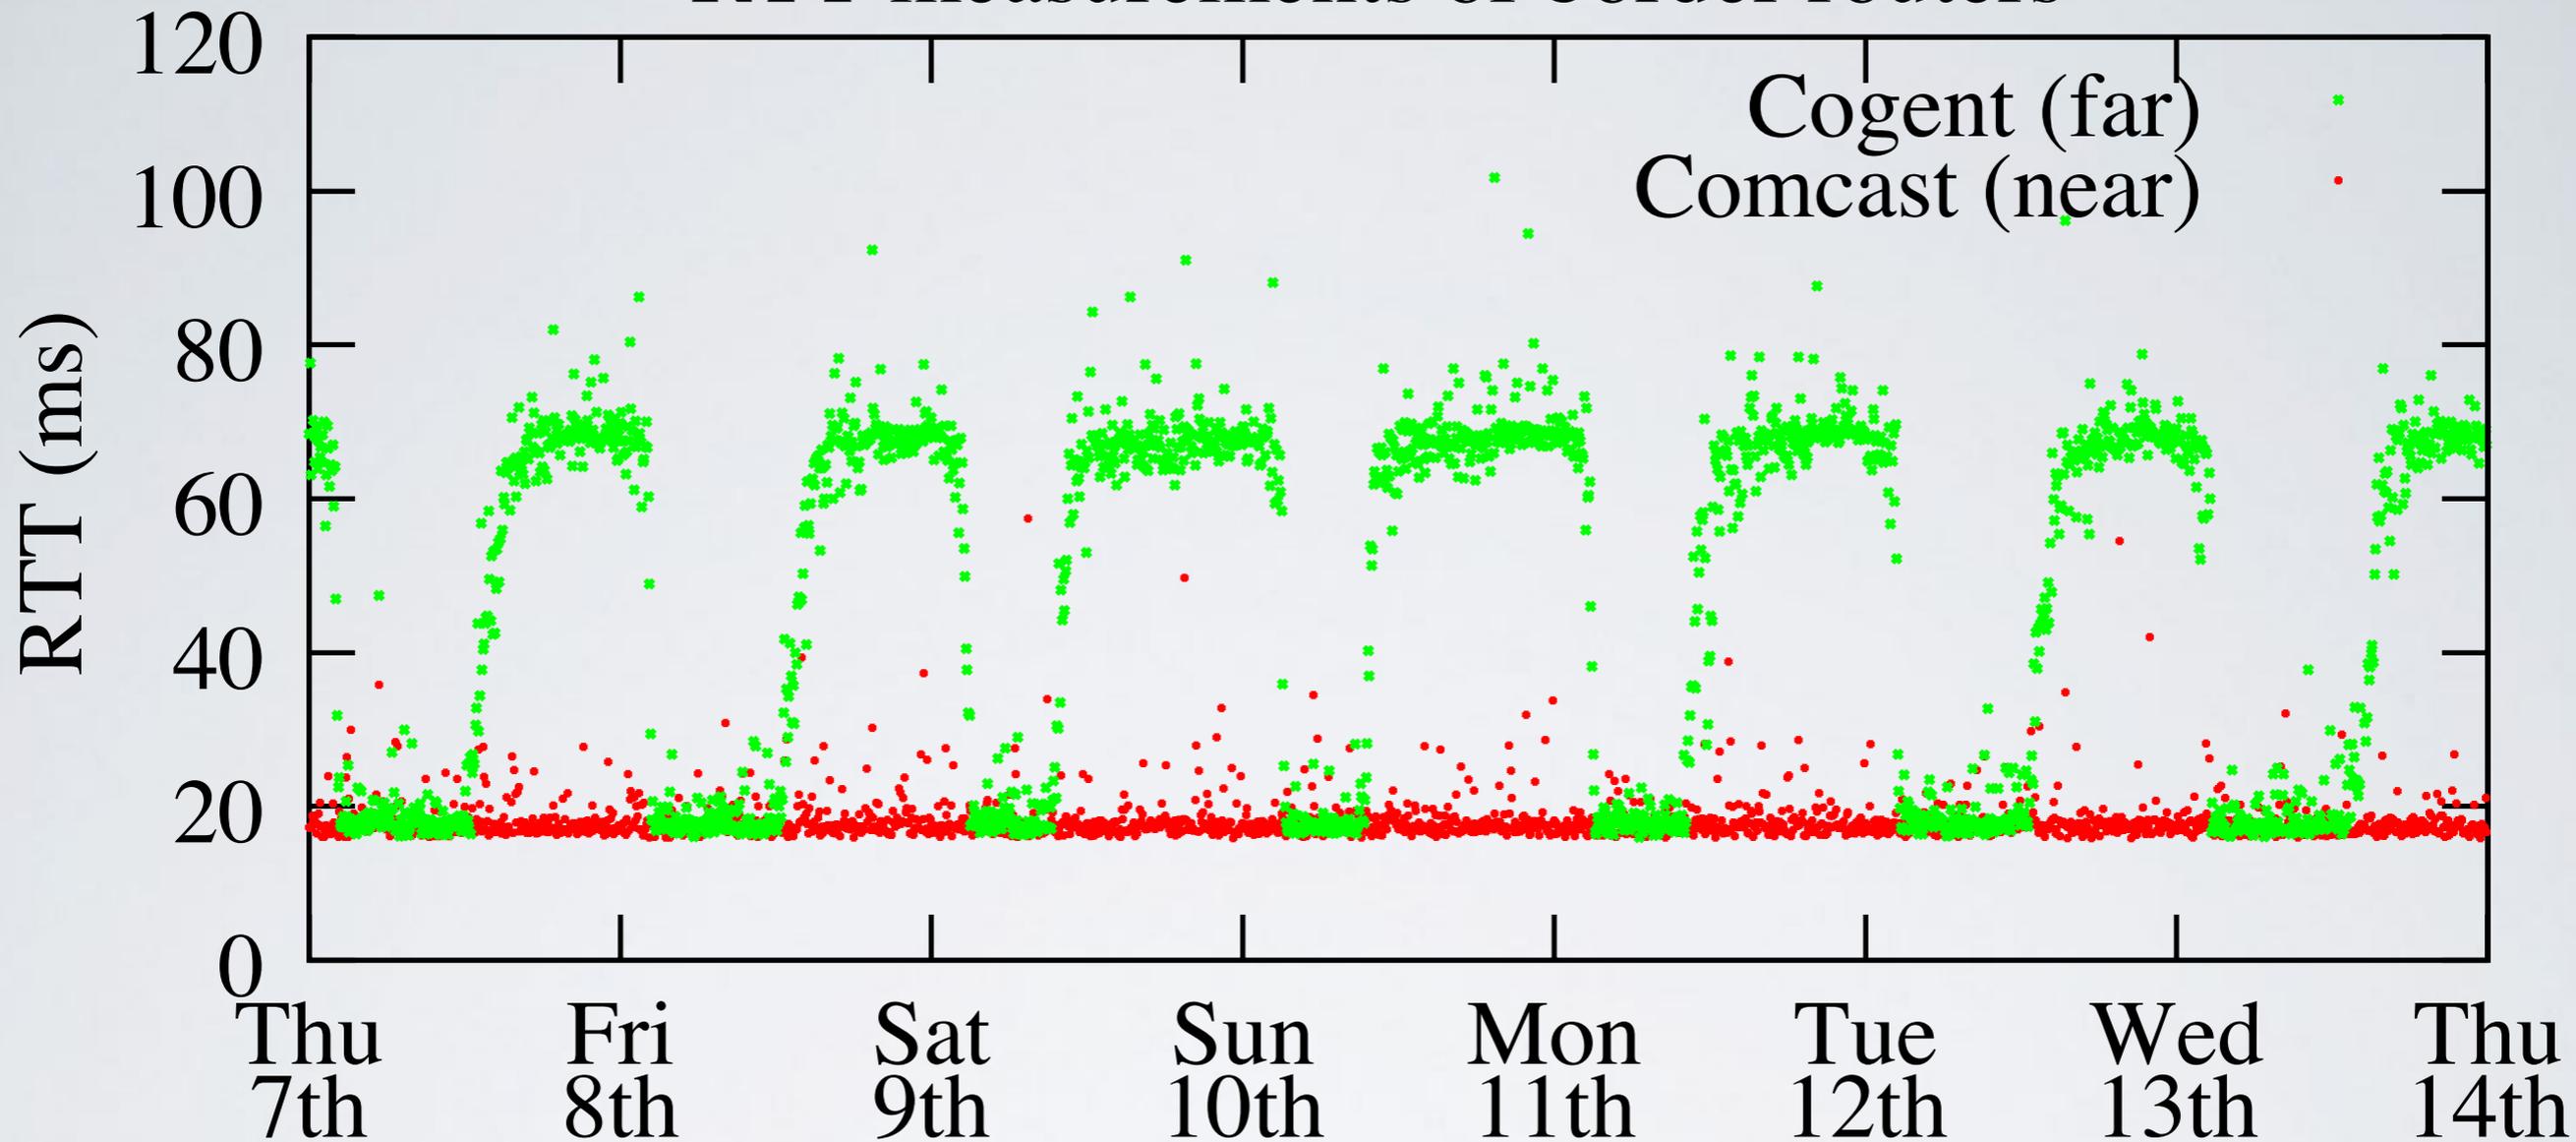
Vantage Point

**Border Routers on
Interesting Link**



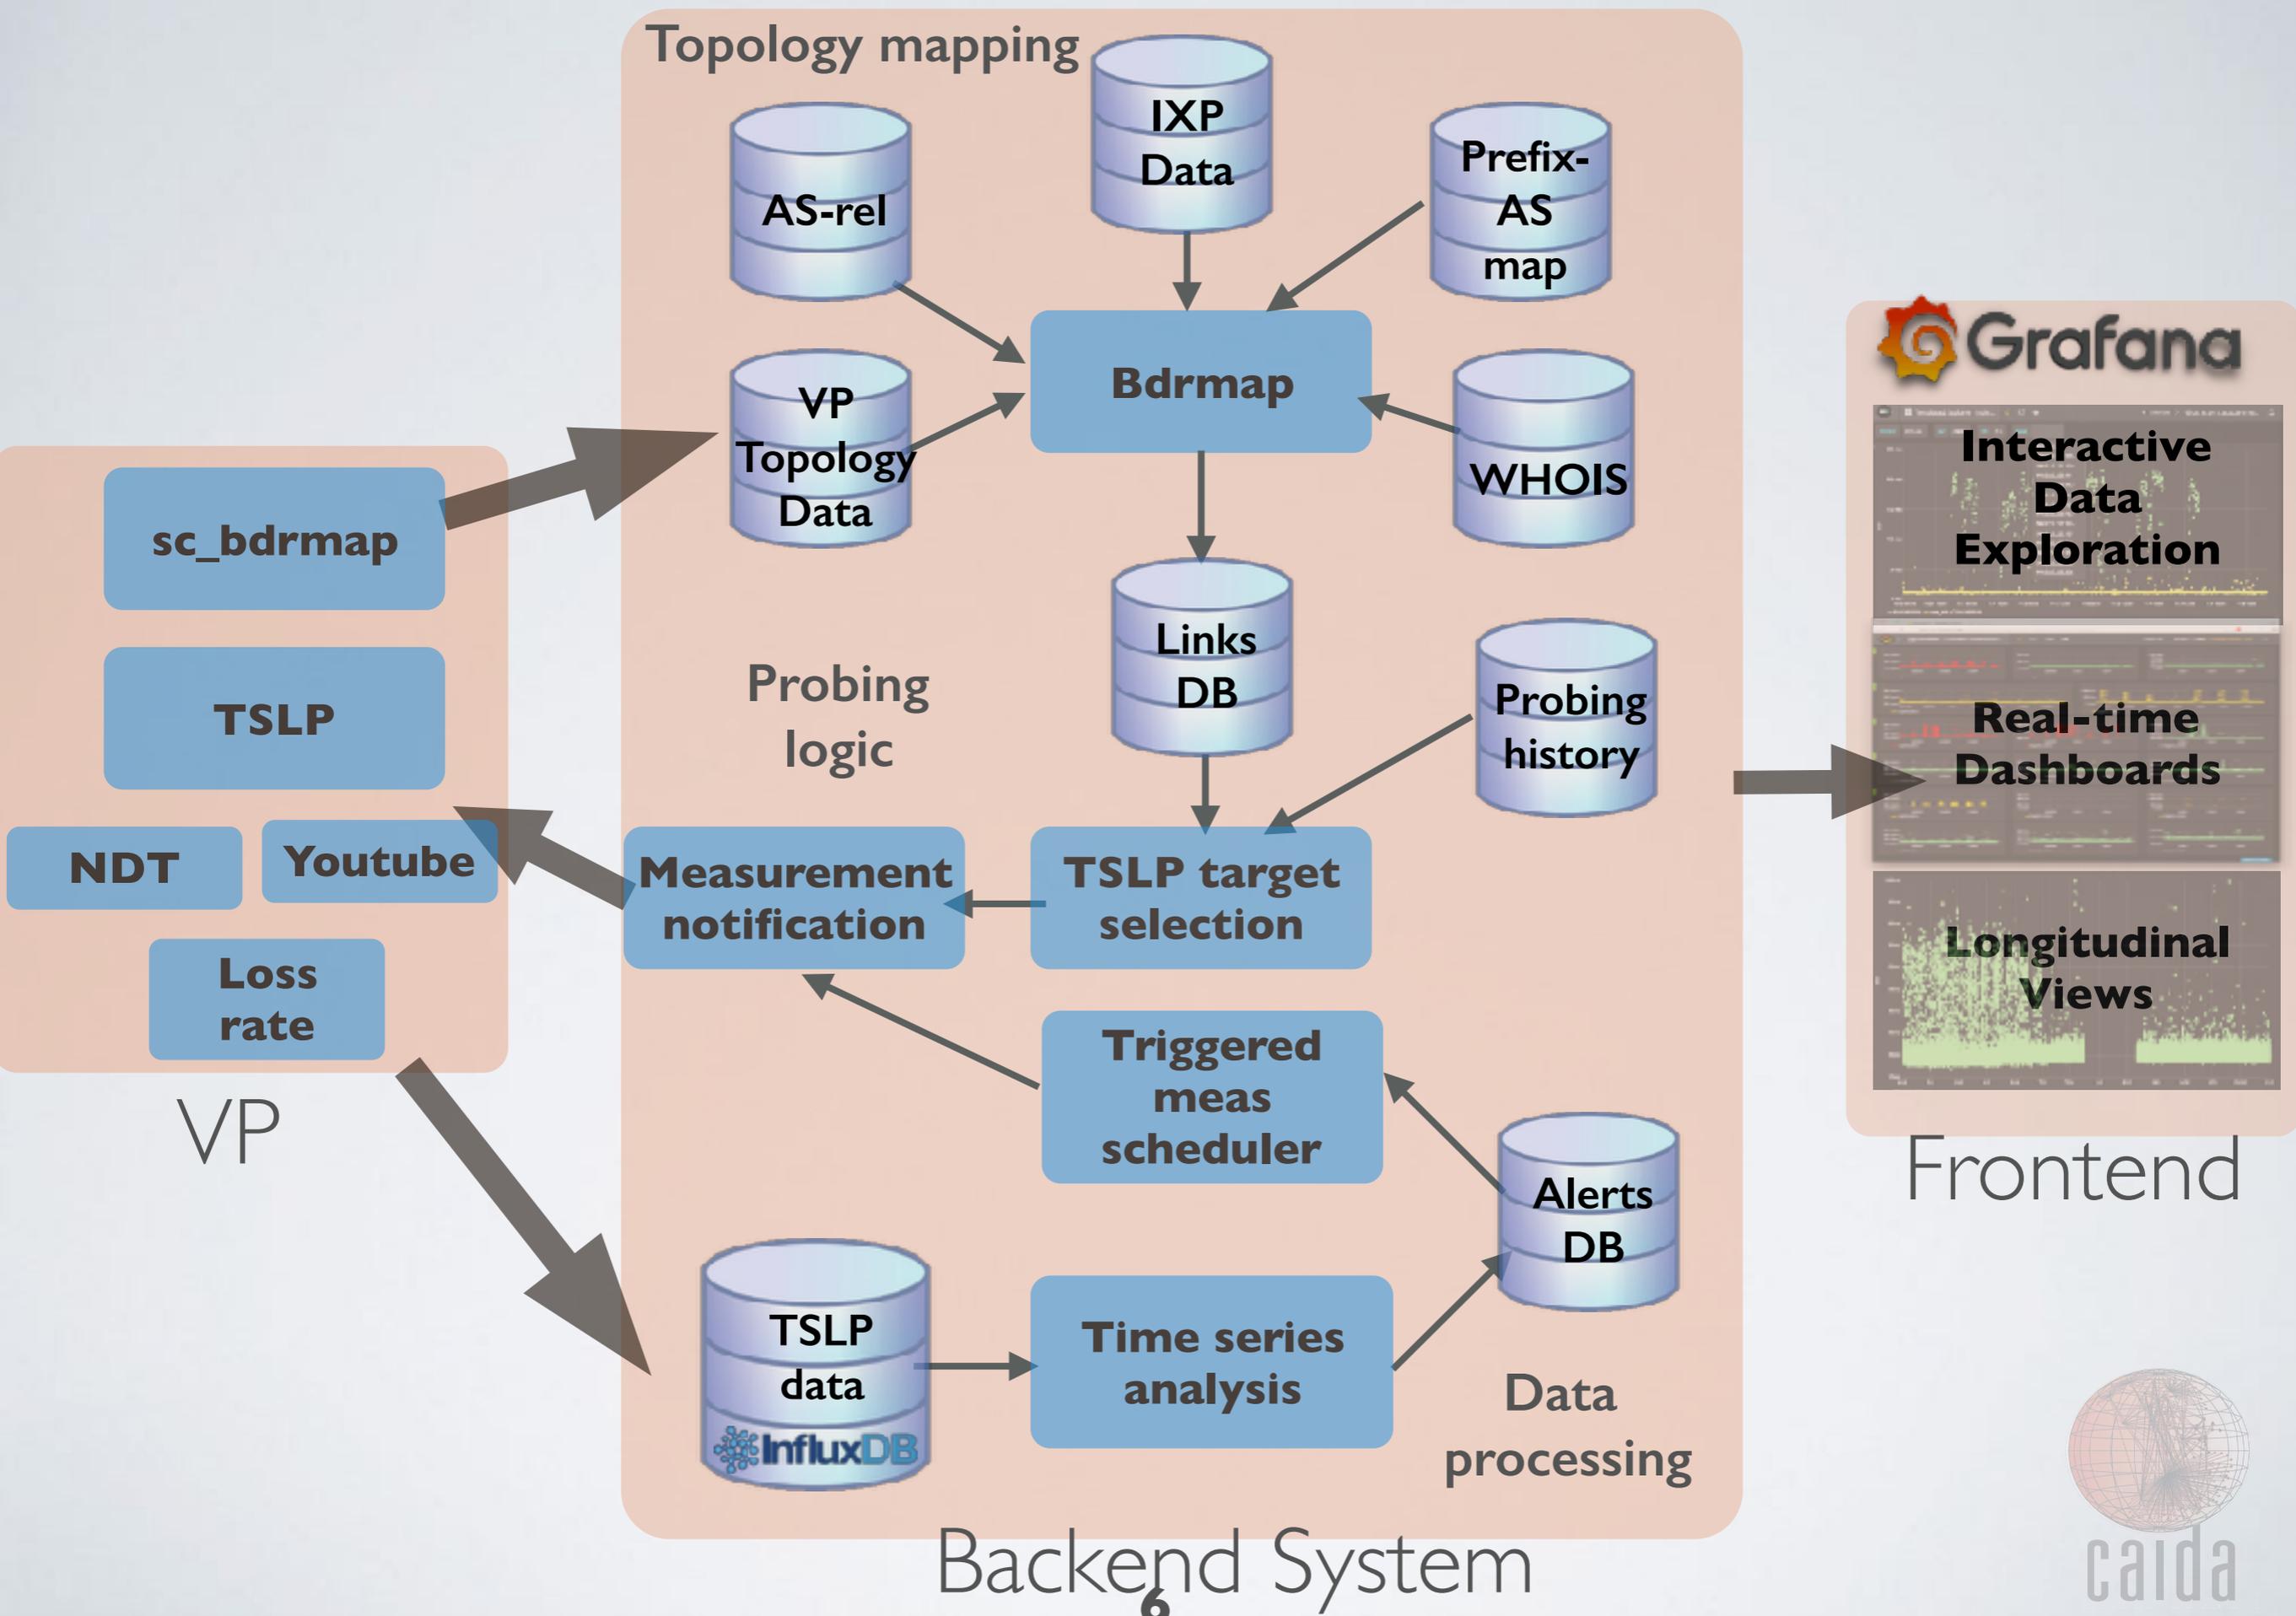
(repeat to obtain a time series)

RTT measurements of border routers



Day of week in November 2013 (in New York)

More congestion on weekend than weekdays. Monday 11th was Veterans Day



Deployment

- bdrmap + TSLP deployed since March 2016 on Ark VPs
- In all: **86 VPs in 47 networks in 24 countries**
- Trial deployment on ~15 Bismark routers

Identifying Congested Links

- Looking for links that are **persistently congested** due to a mismatch between peak-time traffic load and installed capacity
- Intuition: look for periods of elevated latency from a link that correlate across days using an **autocorrelation** function
- (much more complicated than this in practice)

Validation: Operator Feedback

- Validated our inferences with operators from two large U.S. access ISPs
- Asked ISP A about 7 links (all inferred congested) and ISP B about 20 links (10 inferred congested and 10 inferred uncongested)
- **Our inferences were correct in each case**

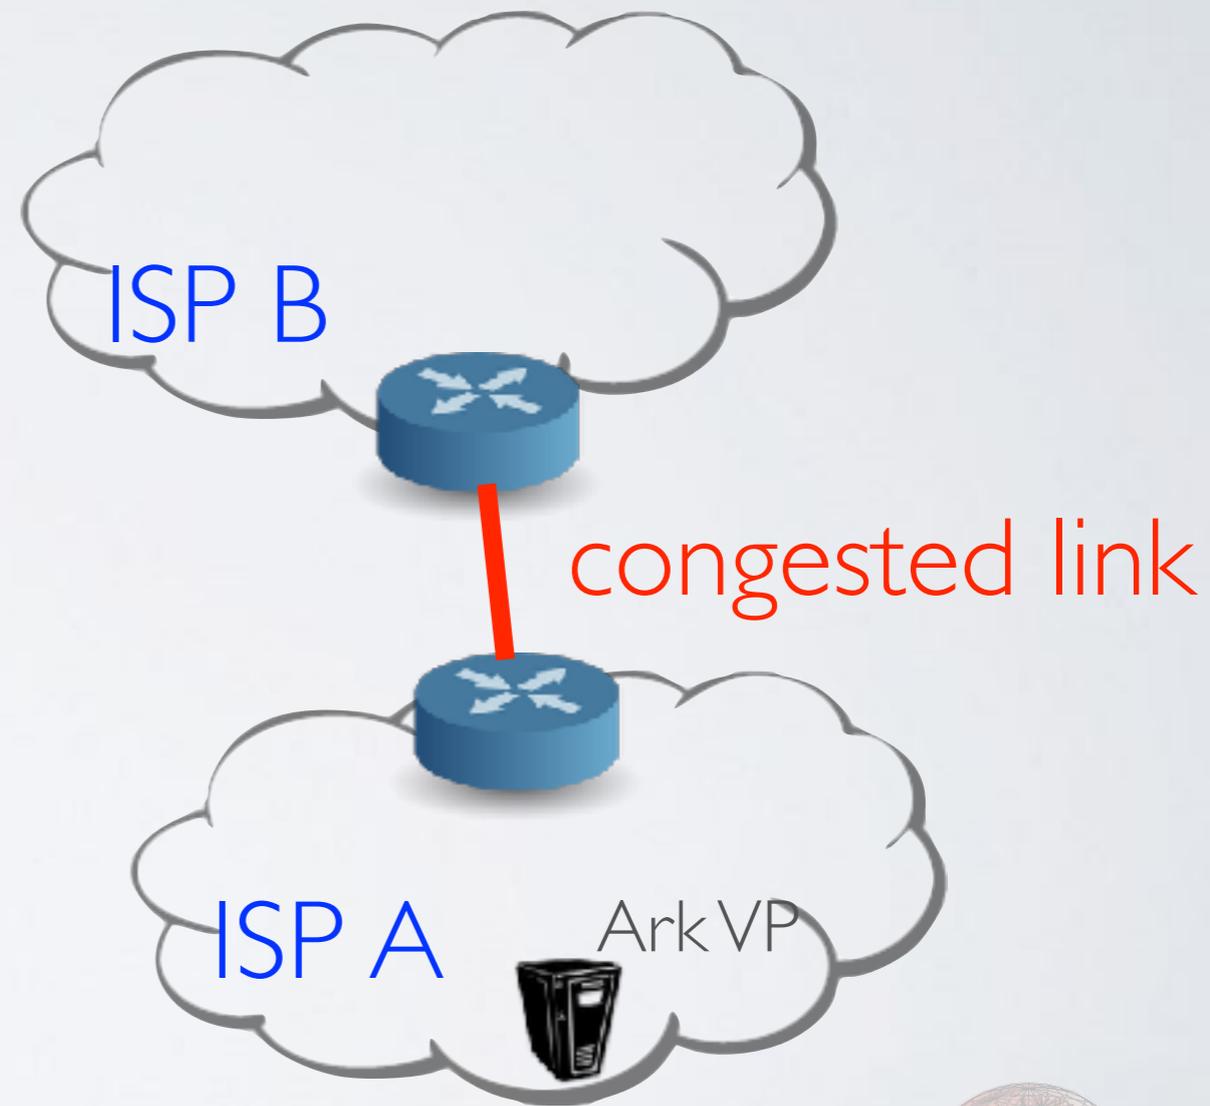
Impact on Throughput

- Periodic NDT tests from a VP to M-lab NDT server behind a previously identified congested link
- Traceroute to NDT server to get forward path
- Analyze M-lab traceroutes to get reverse path
- Network Diagnostic Test (NDT):
Measures upstream and downstream throughput



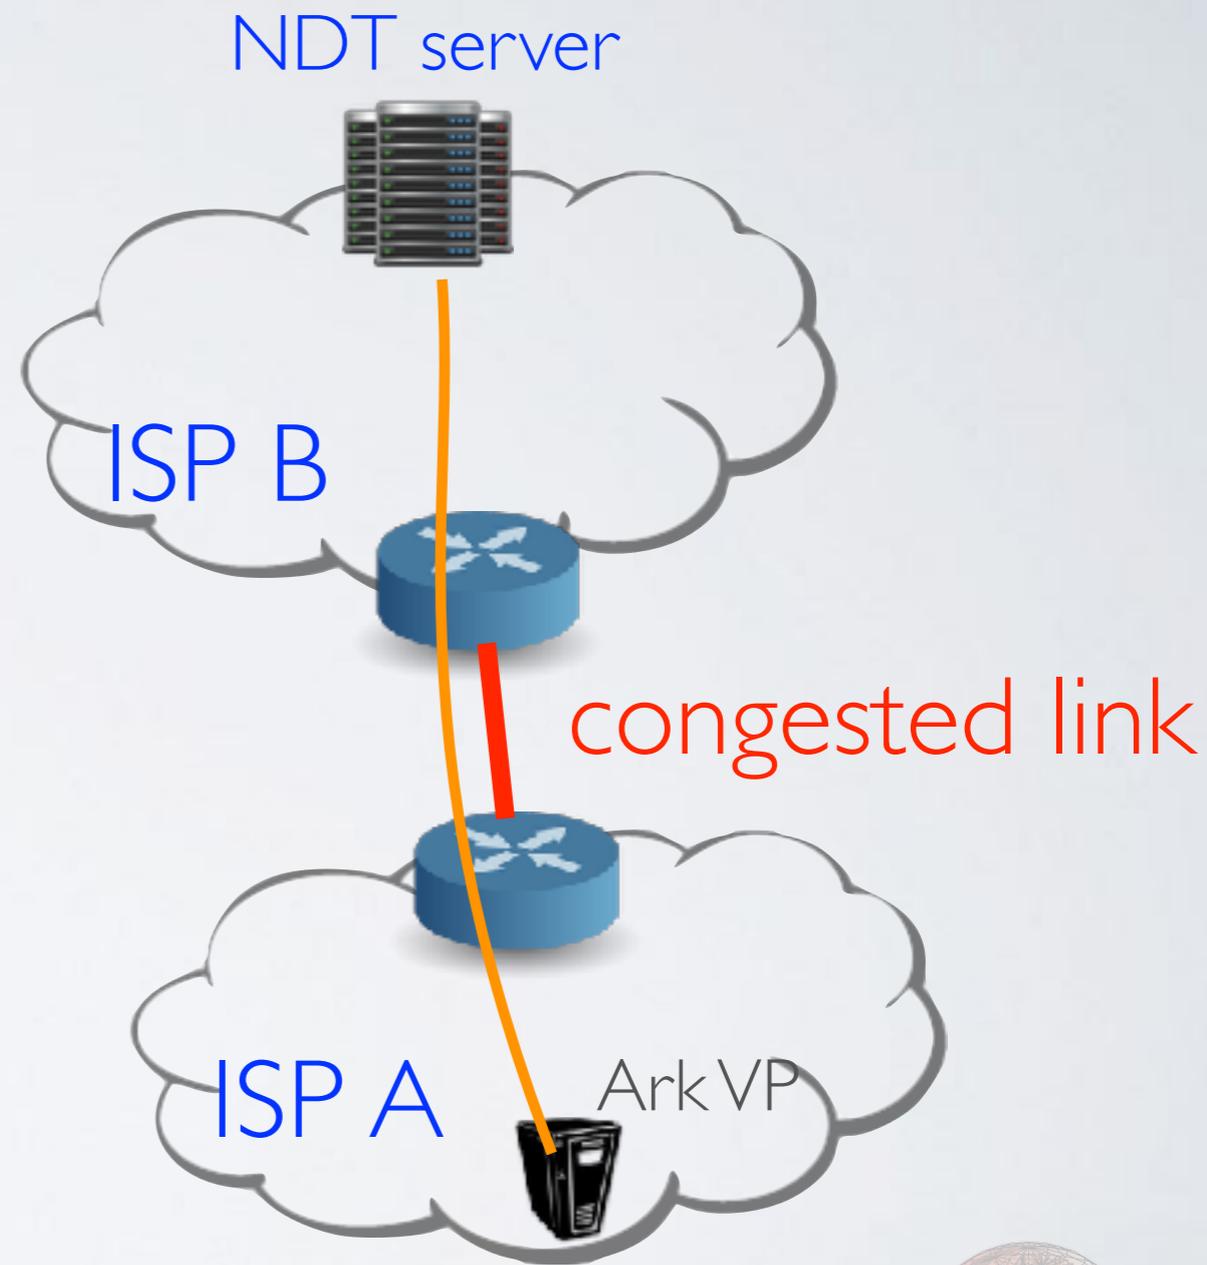
Impact on Throughput

- Periodic NDT tests from a VP to M-lab NDT server behind a previously identified congested link
- Traceroute to NDT server to get forward path
- Analyze M-lab traceroutes to get reverse path
- Network Diagnostic Test (NDT):
Measures upstream and downstream throughput

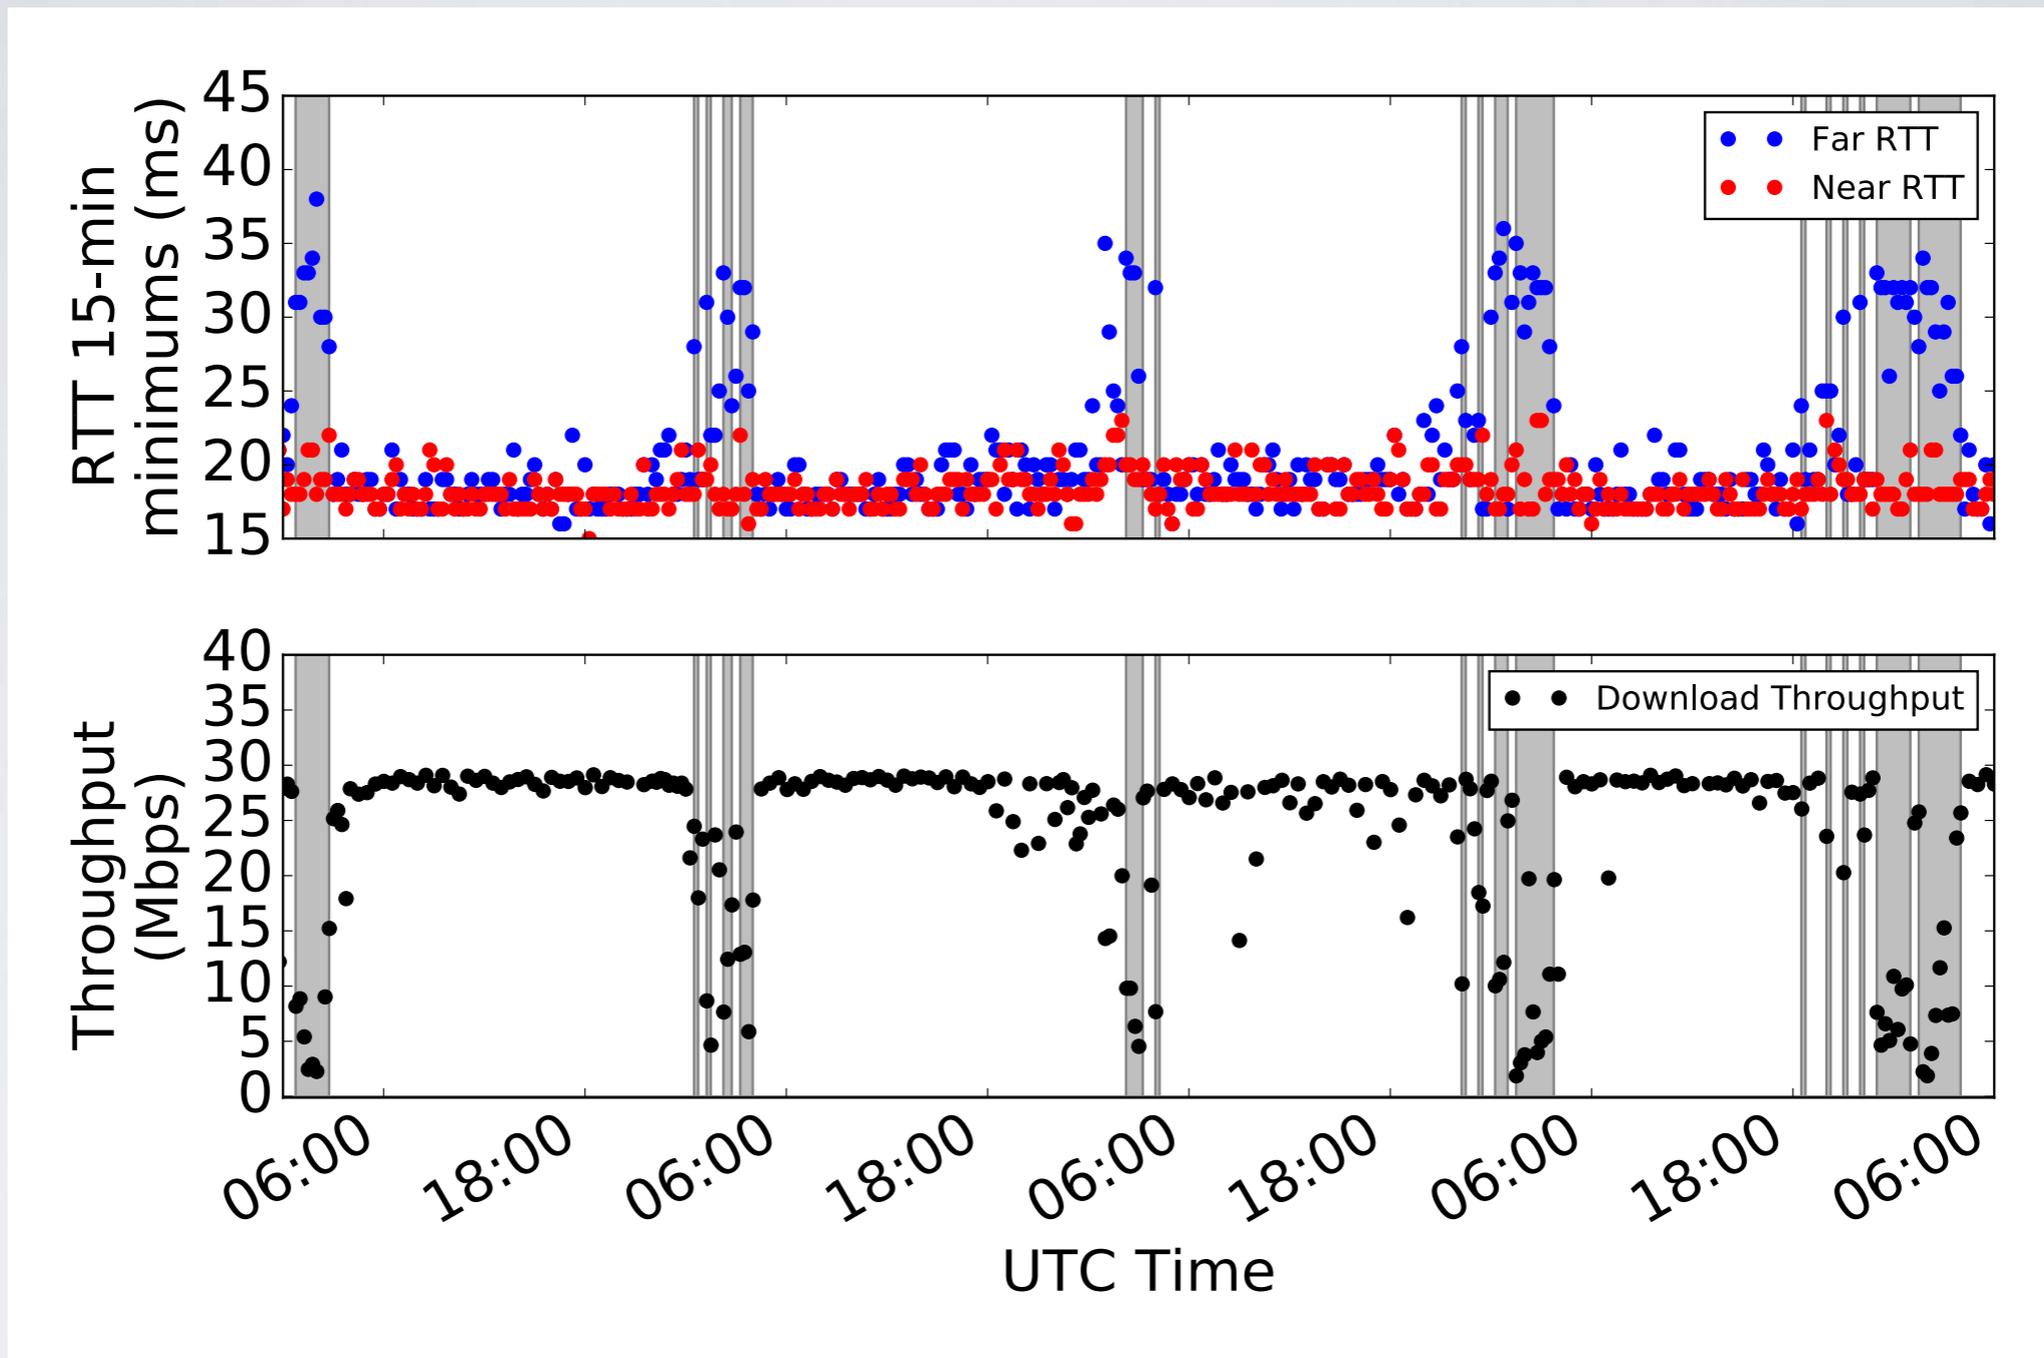


Impact on Throughput

- Periodic NDT tests from a VP to M-lab NDT server behind a previously identified congested link
- Traceroute to NDT server to get forward path
- Analyze M-lab traceroutes to get reverse path
- Network Diagnostic Test (NDT):
Measures upstream and downstream throughput



Impact on Throughput



- Throughput shows negative correlation with TSLP latency: **Average throughput ~27Mbps uncongested, ~8Mbps congested**

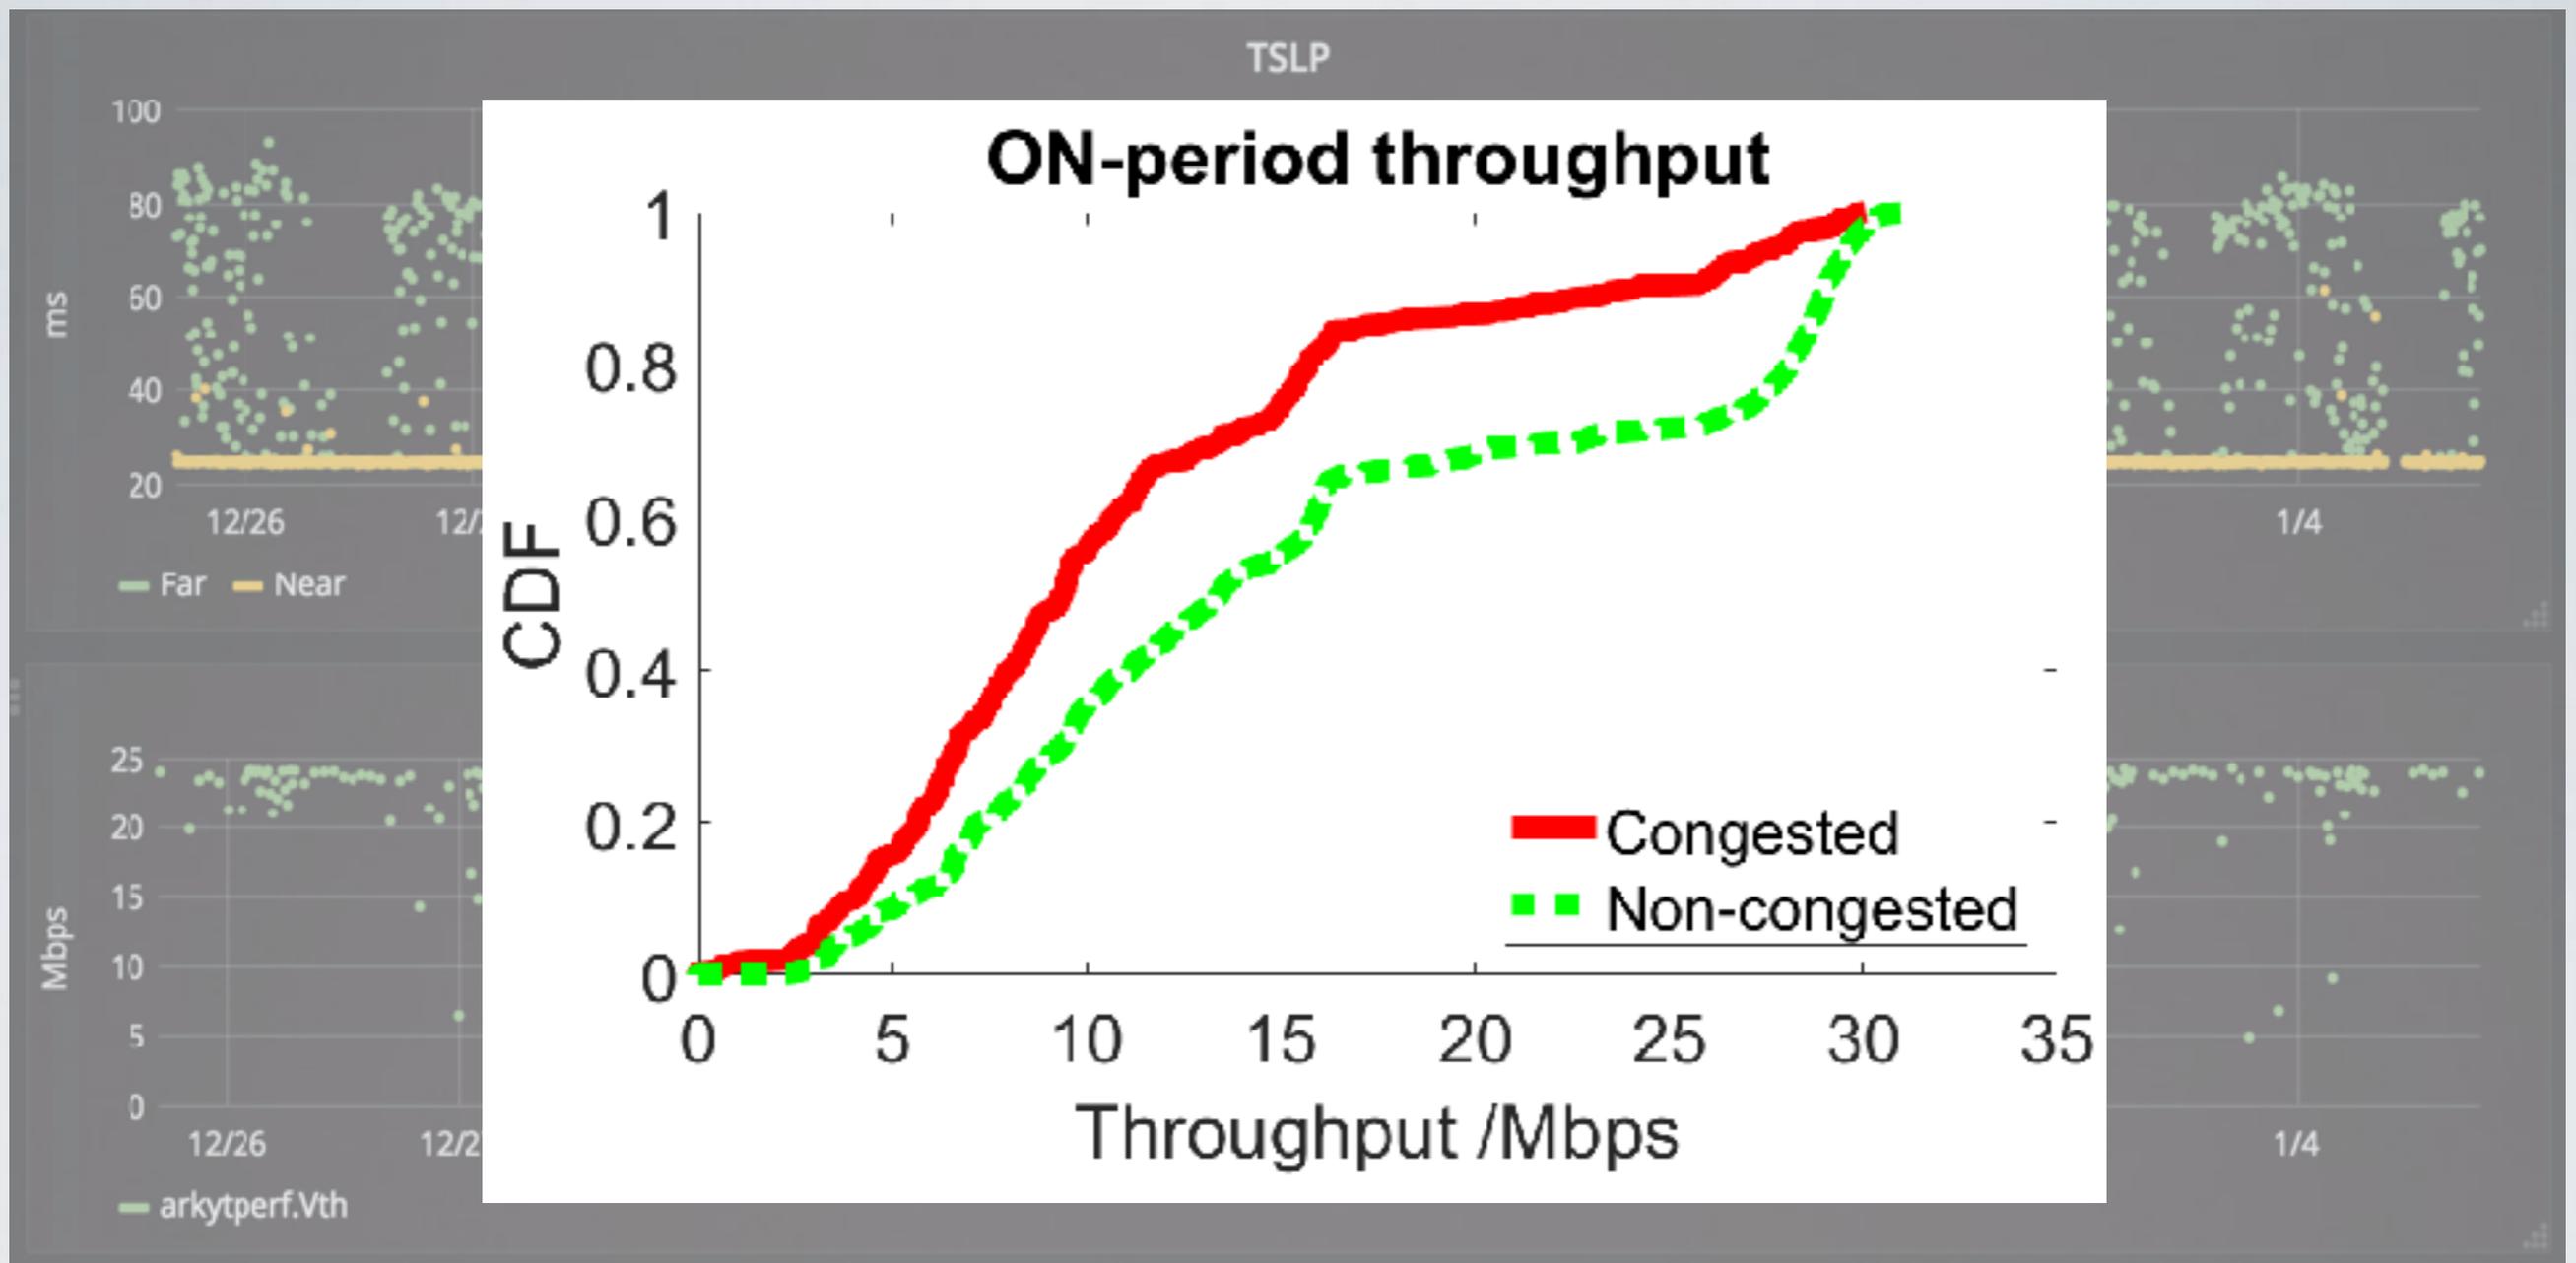
Youtube Streaming

- Do our inferences of interdomain congestion correlate with degraded Youtube streaming metrics?
- Custom YouTube test from Ark and SamKnows VPs measures throughput, startup delay, rebuffering etc.
 - 6 Ark VPs and 12 SamKnows VPs
- Traceroute to the YouTube CDN server to identify interdomain link in forward path

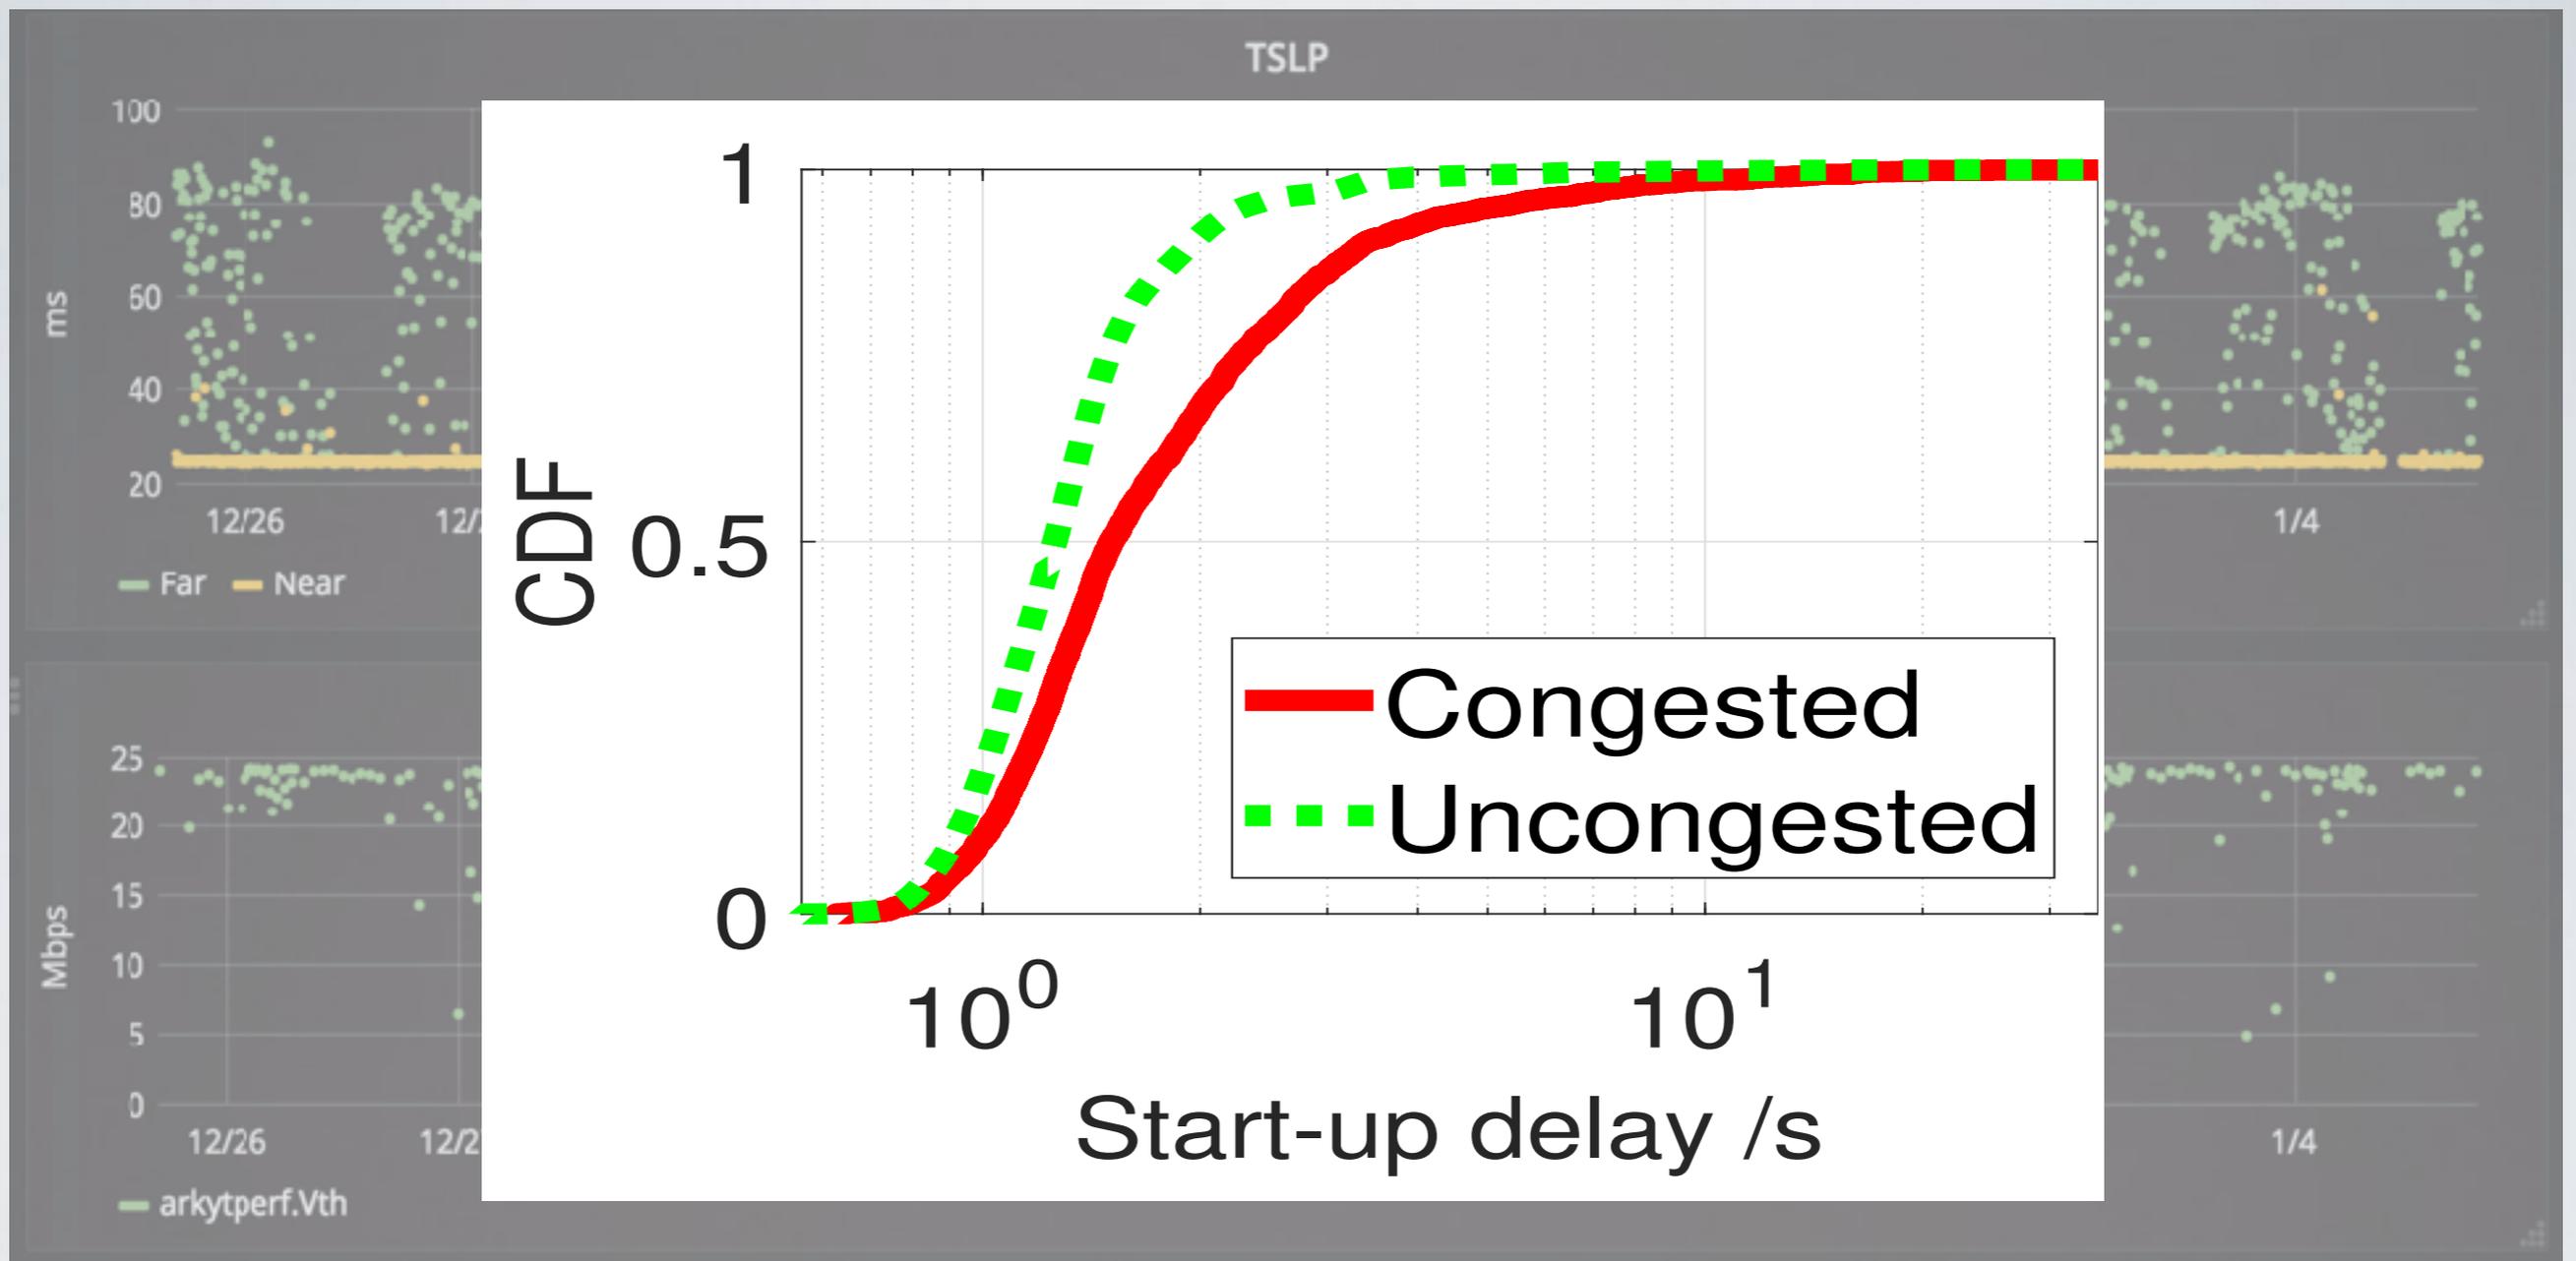
Impact on User QoE (Youtube)



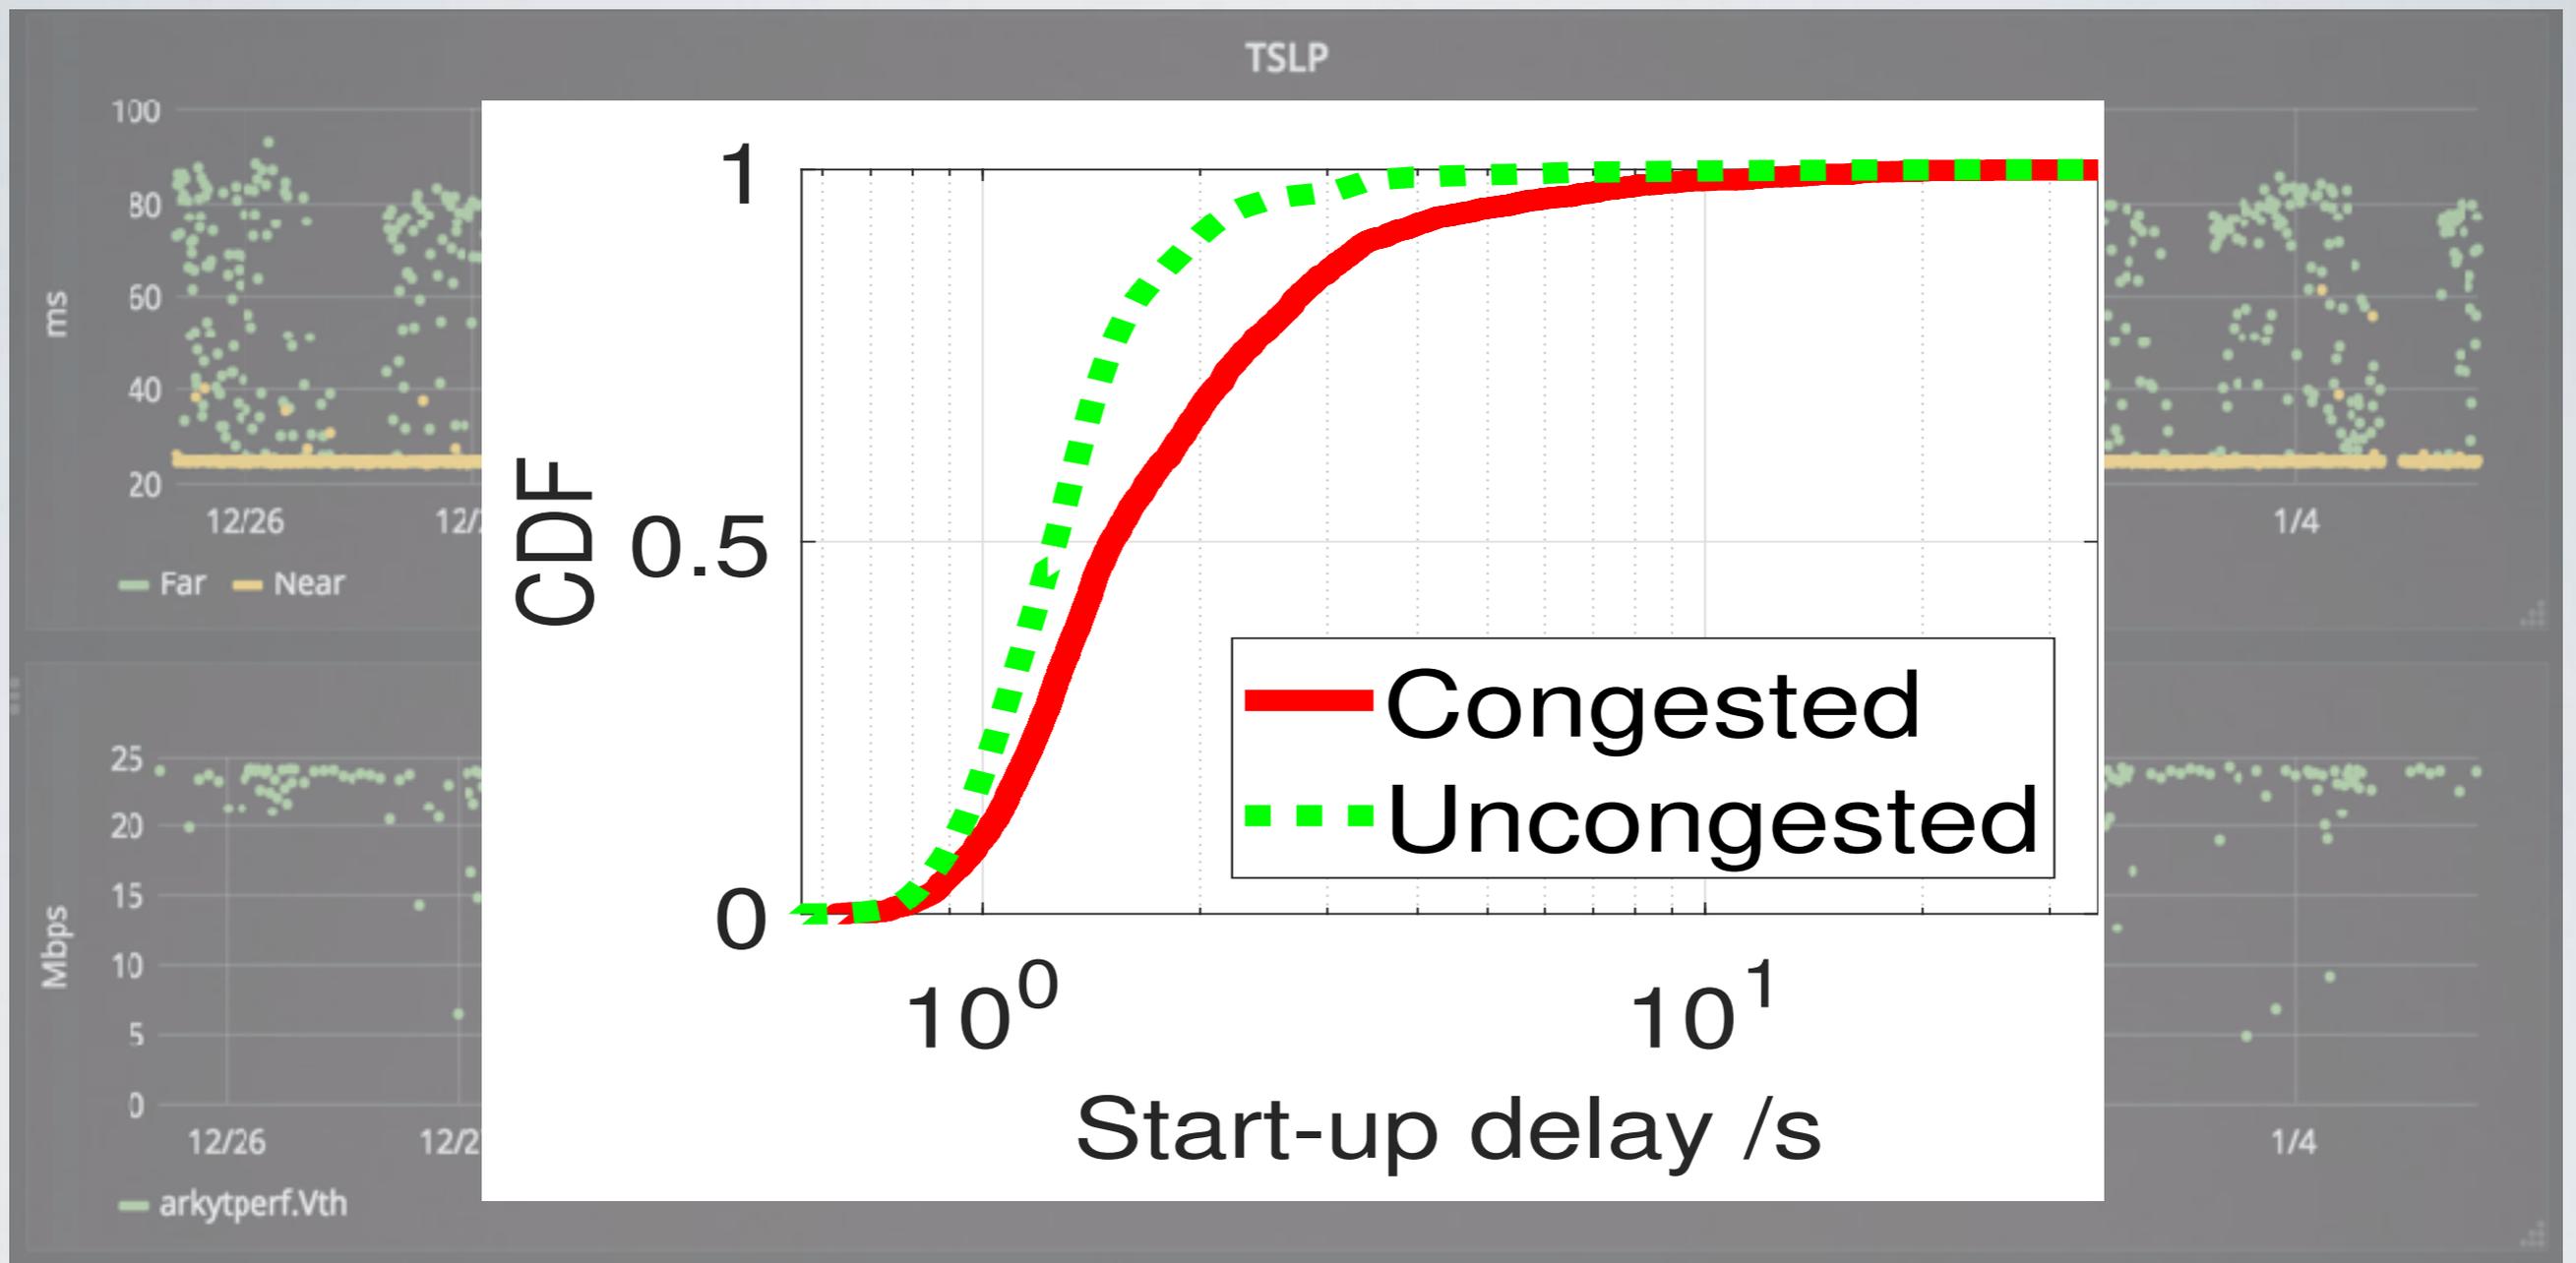
Impact on User QoE (Youtube)



Impact on User QoE (Youtube)



Impact on User QoE (Youtube)



Measurable impact on ON throughput, startup delay, and test failure rate

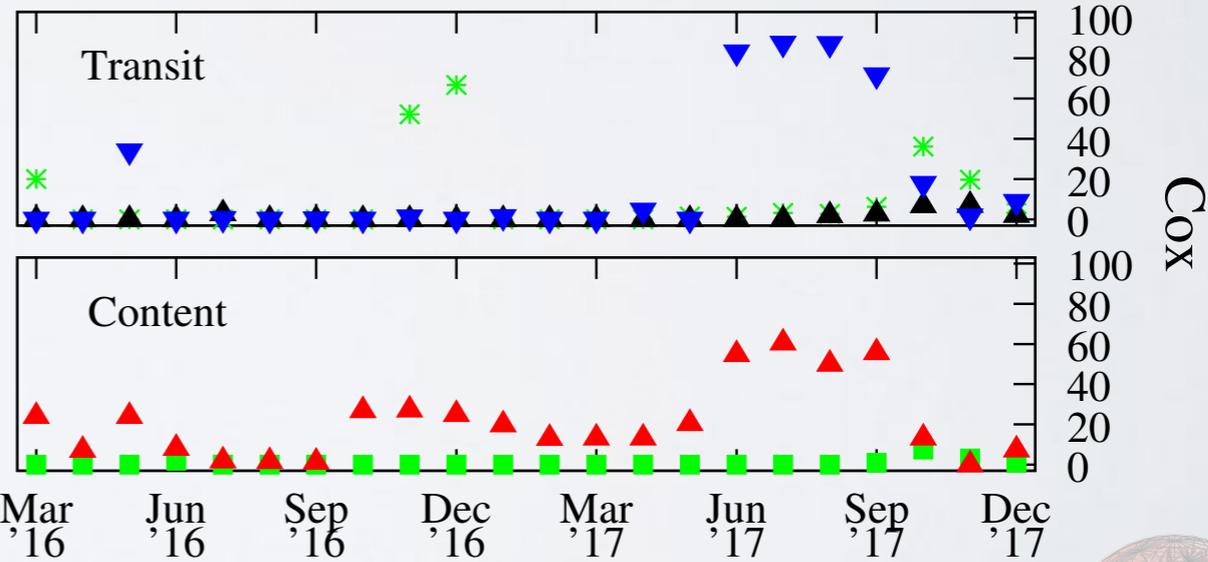
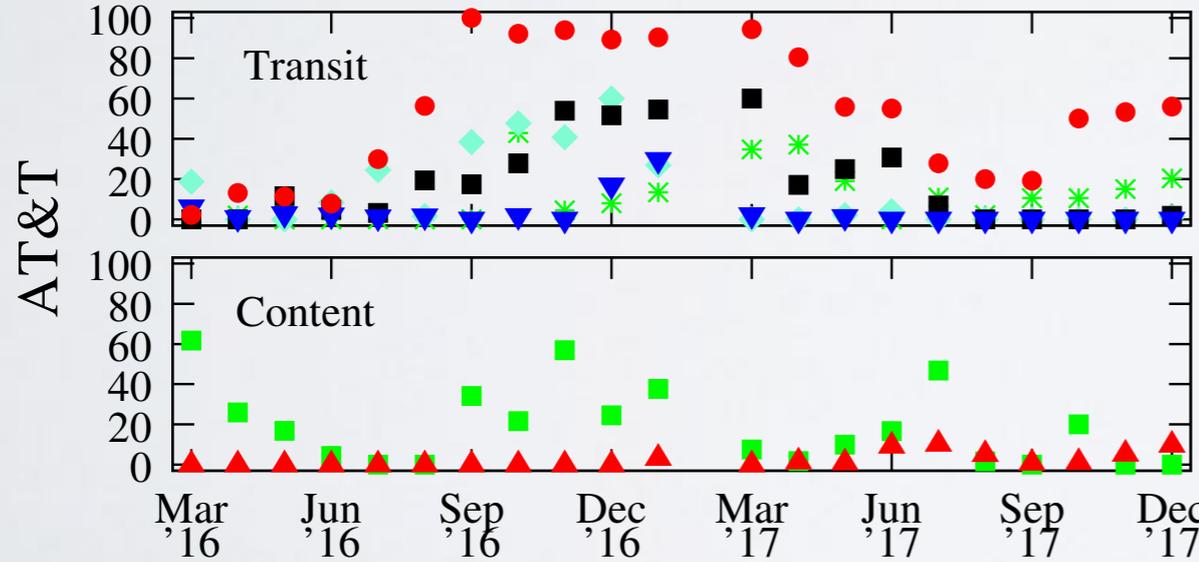
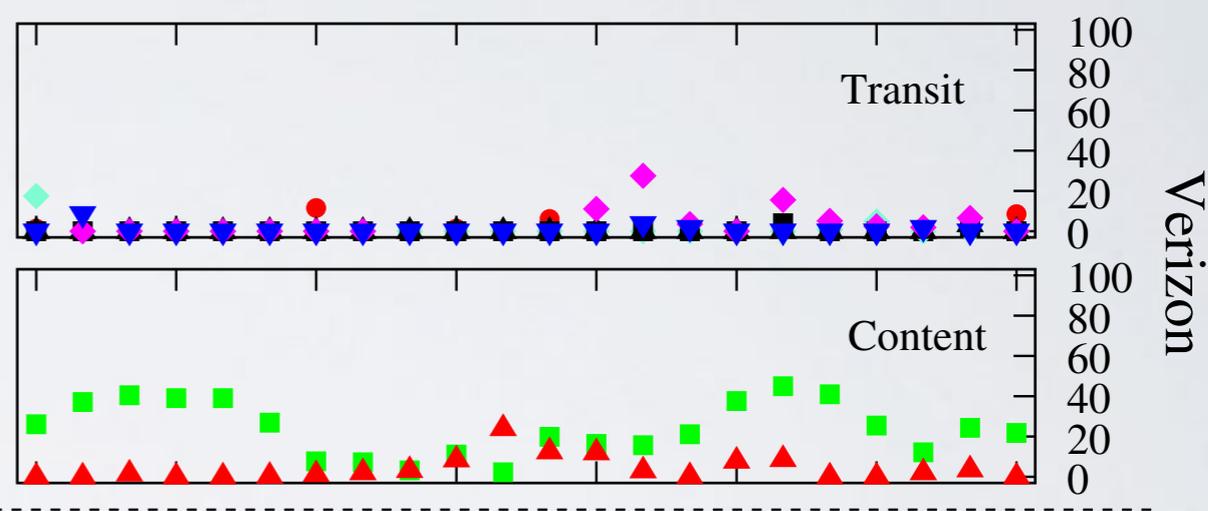
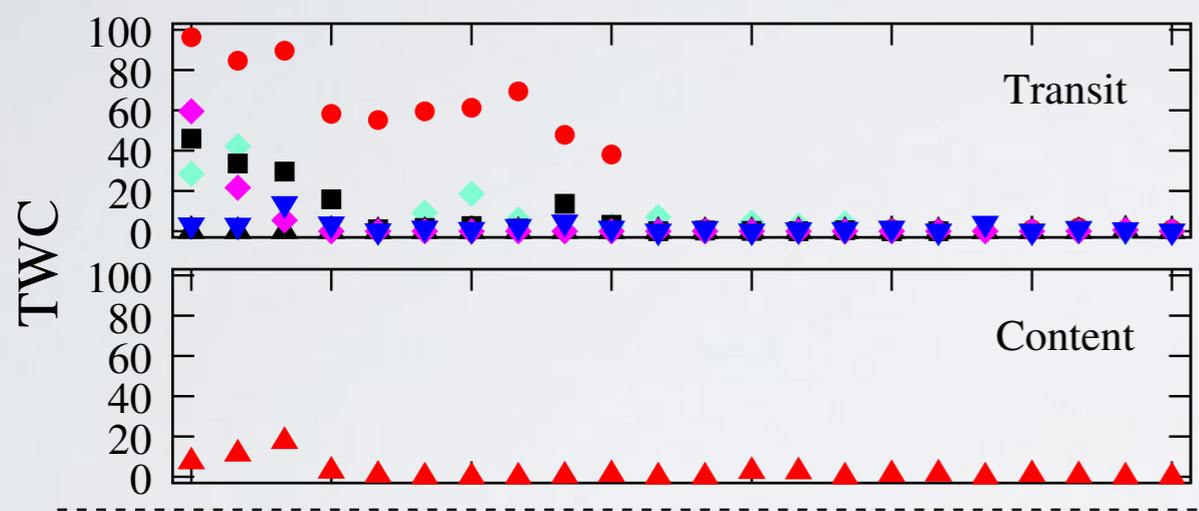
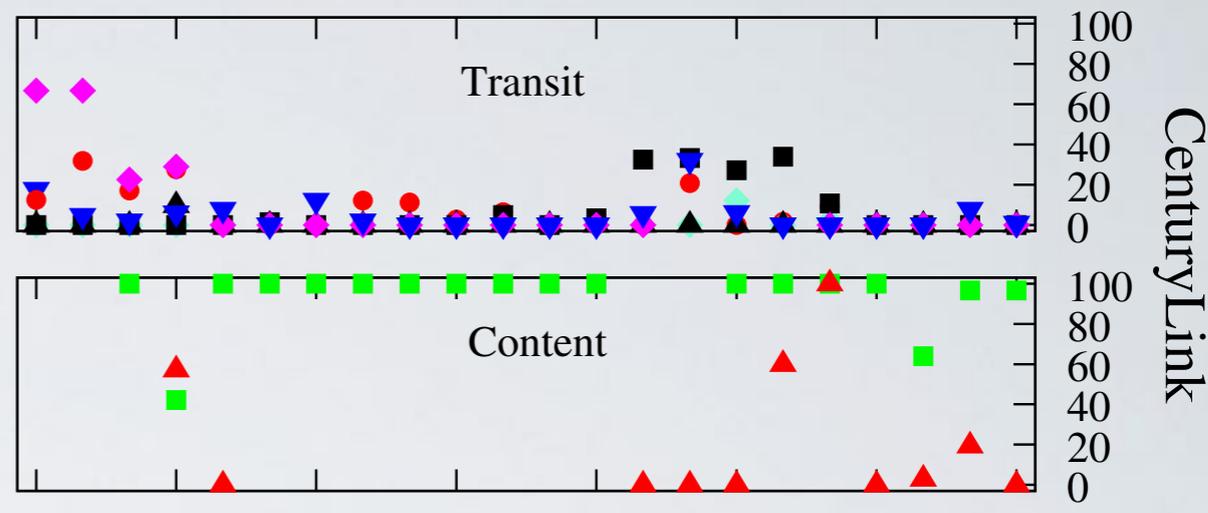
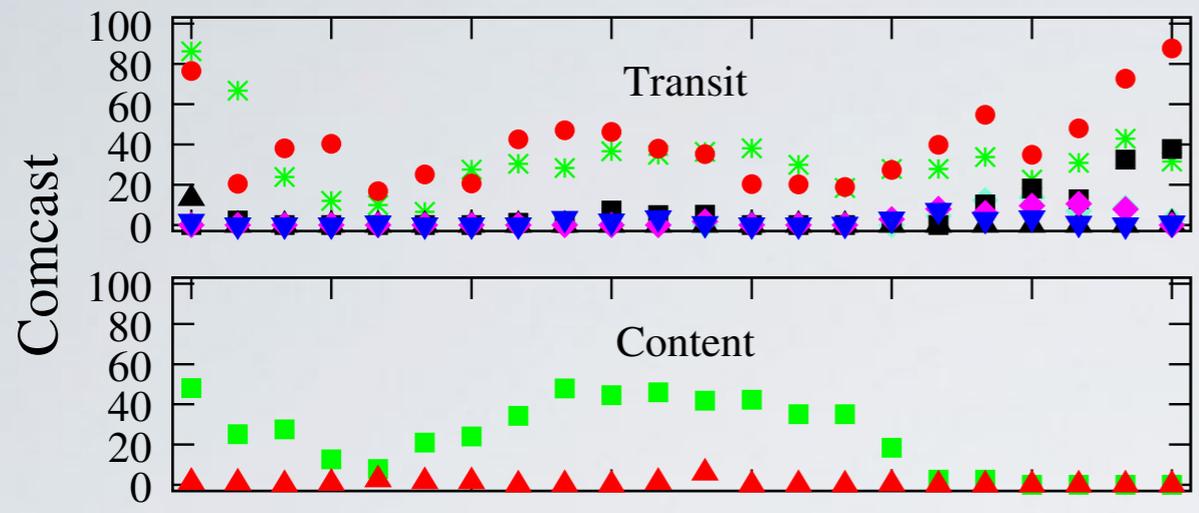
Longitudinal Study

- Focused on interdomain links of 8 large access ISPs in the U.S. to transit providers and peers
- TSLP data collection from March 2016 to present
- Which access providers have the most congested links?
Which transit/content providers are most often congested to access providers? Do we see trends over time?



Summary of findings

- **No evidence of widespread congestion**
- Small fraction of peers of every access provider ever showed evidence of congestion
- We monitored every visible link on every day of our measurement study. Small fraction of these “day-links” (< 10%) for each ISP were inferred congested
- Some transit providers (e.g., TATA) and content providers (e.g. Google) were most often congested



Interactive Visualization, Public Access to Data, Correlation with Other Data

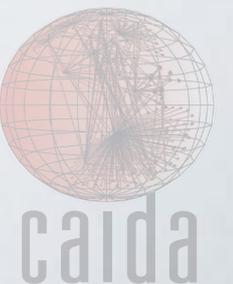
(PANDA!)



Interactive Visualization, Public Access to Data, Correlation with Other Data

(PANDA!)

- “This is a 3-year NSF project, so in year 4 we’ll start thinking of building an interactive visualization system for congestion data that we can make public”
 - Anonymous, AIMS 2015

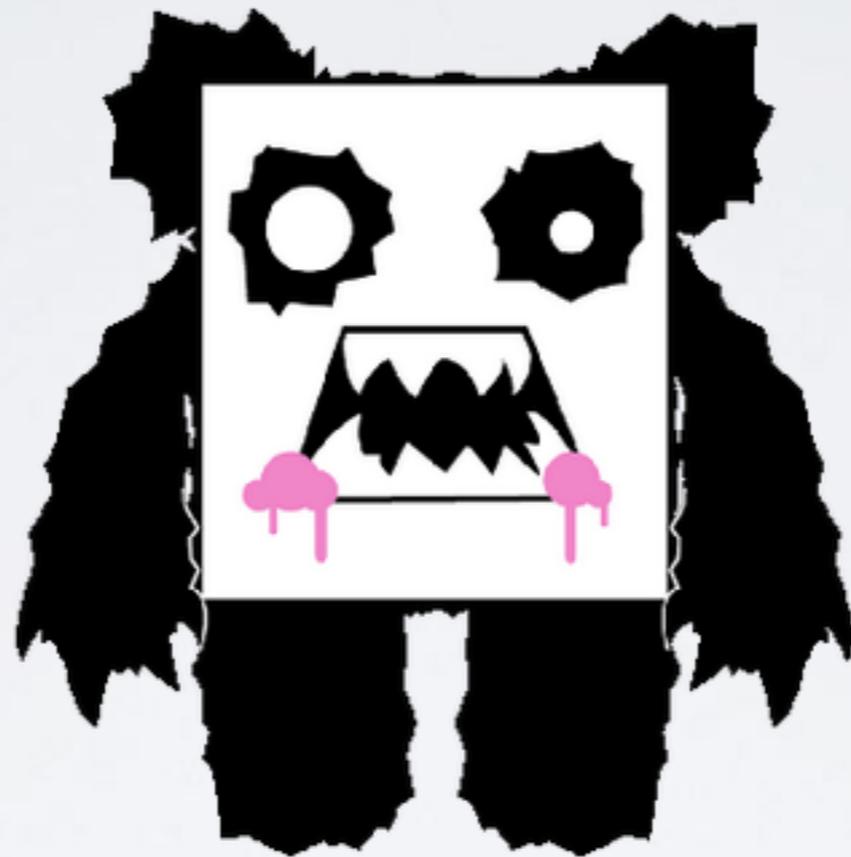


Interactive Visualization and Public Access to Data

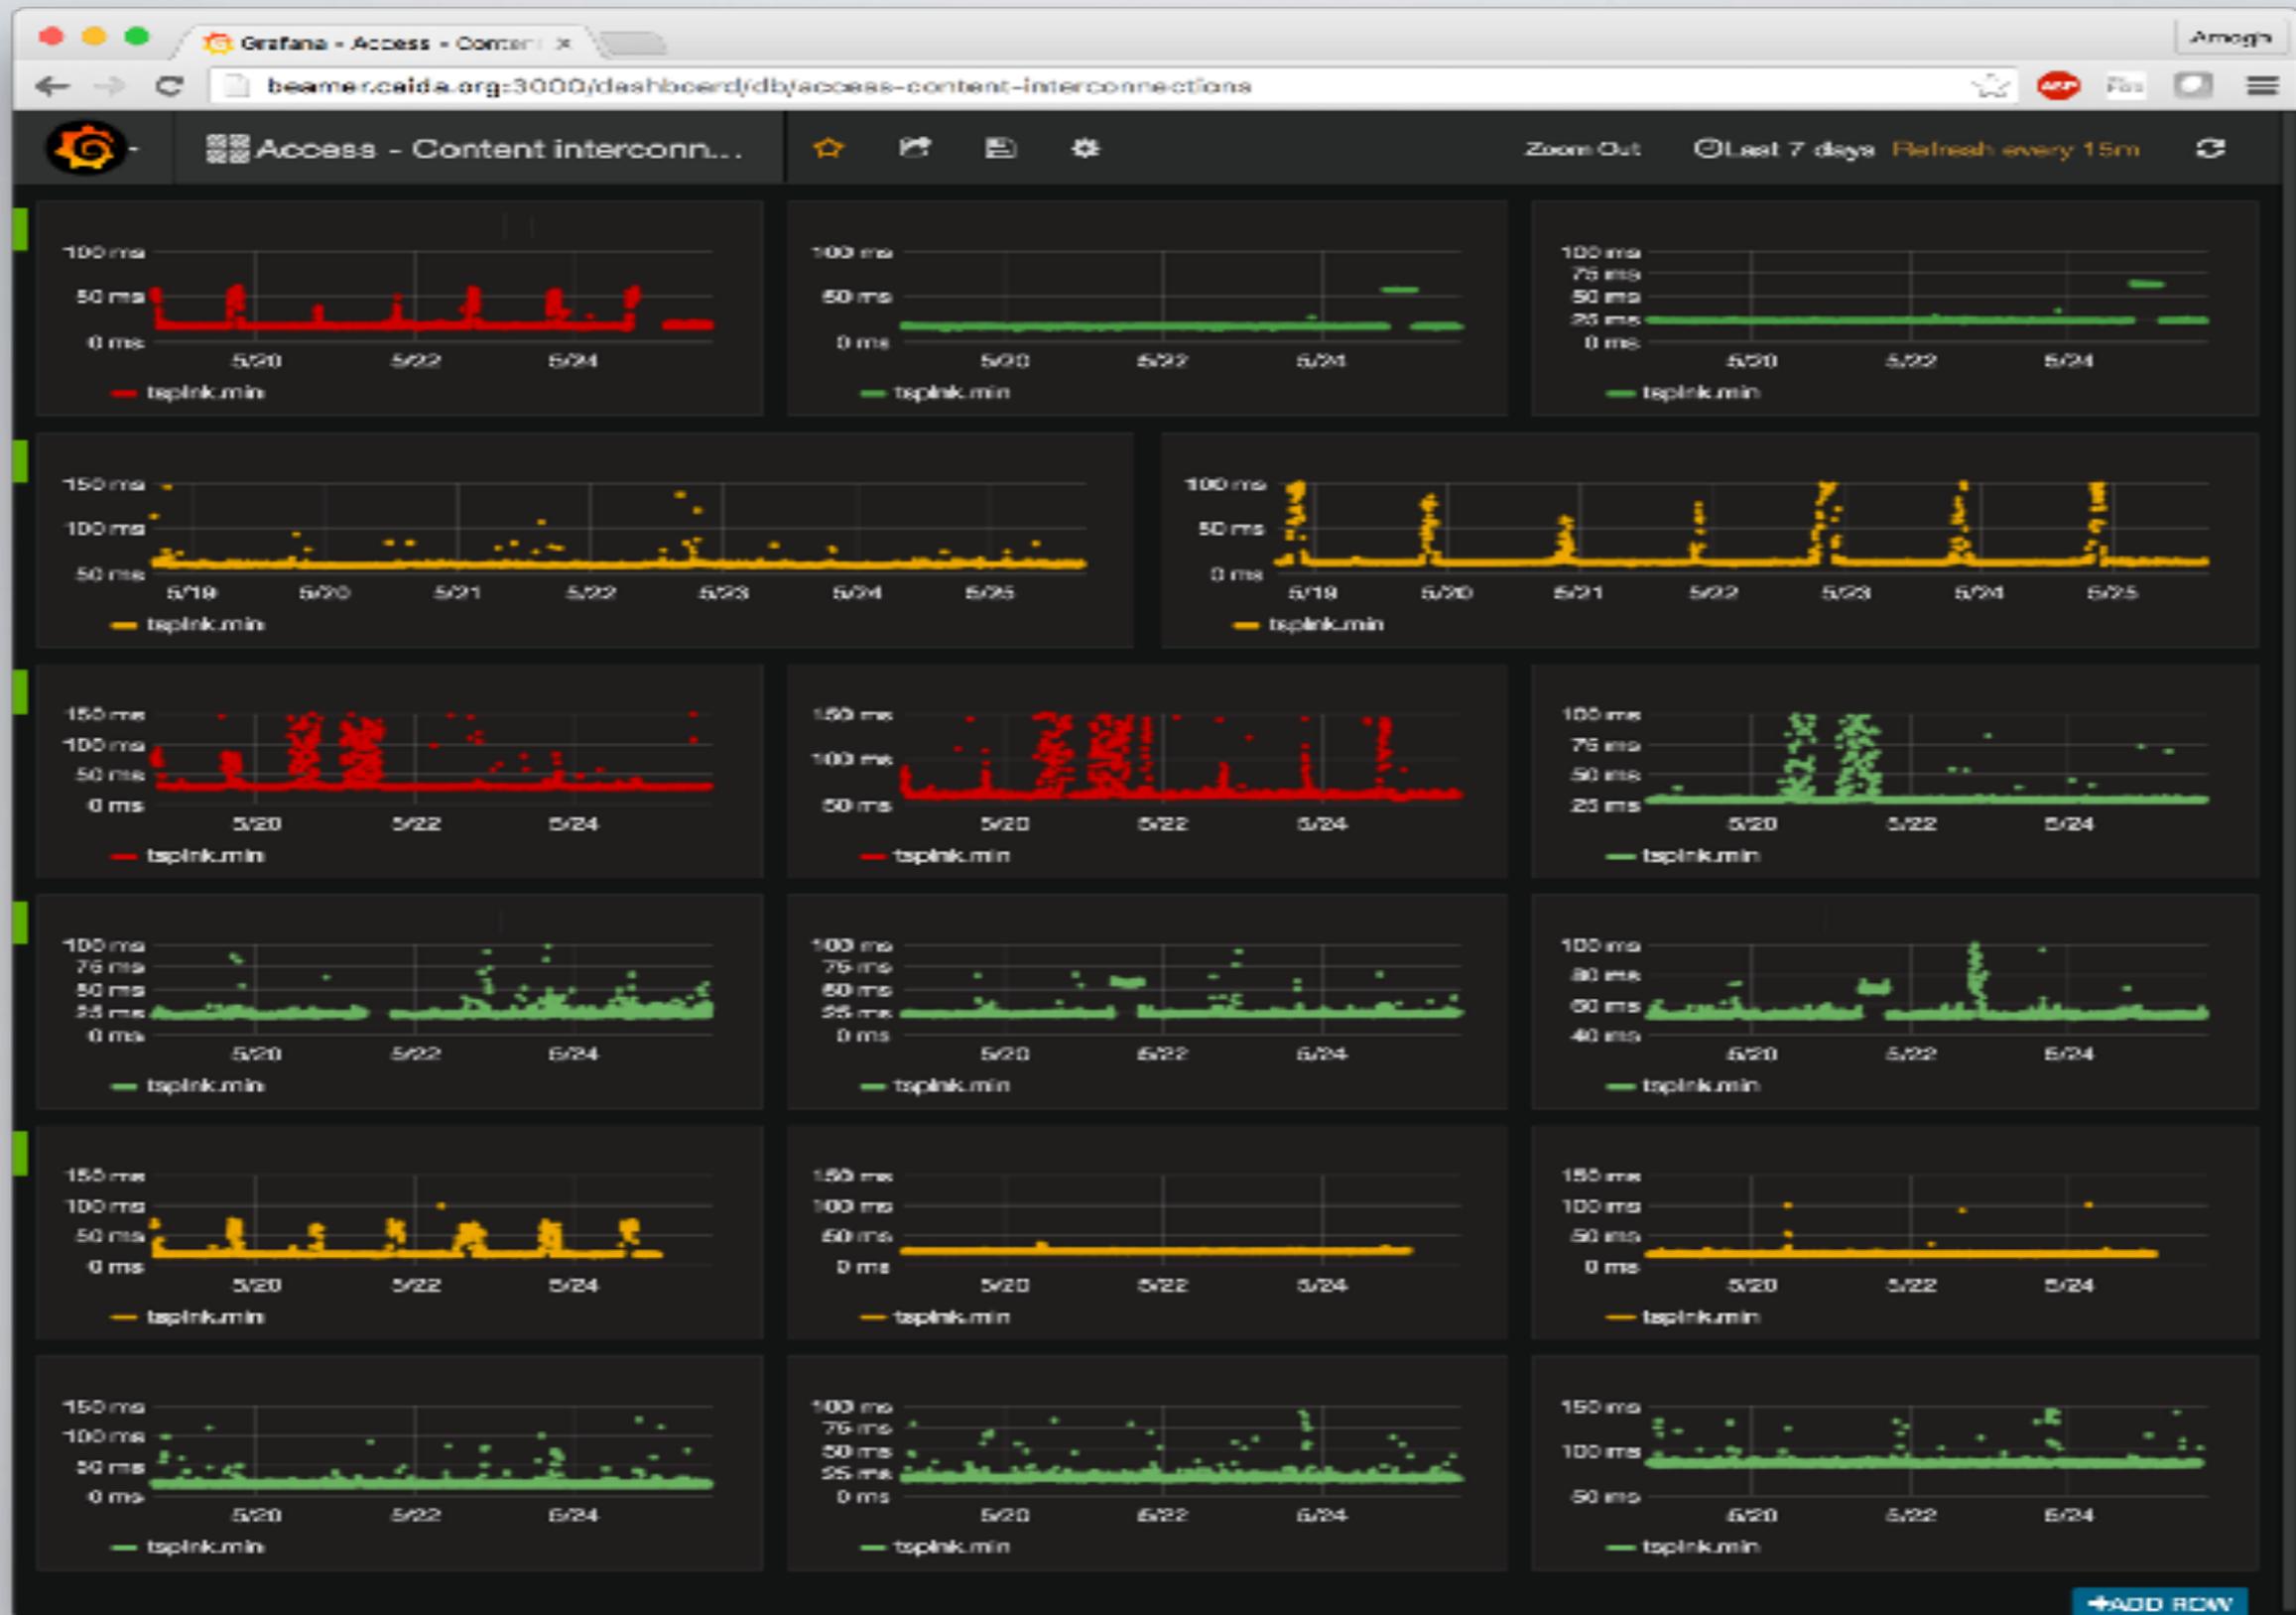
- Anonymous is a pessimist
- We had an interactive visualization system by year 2
- But not open to the public
- Our viz system uses an open-source web framework (Grafana) with influxDB as the underlying data store
- **Coming soon: going public (aka crowdsourcing breakage)**



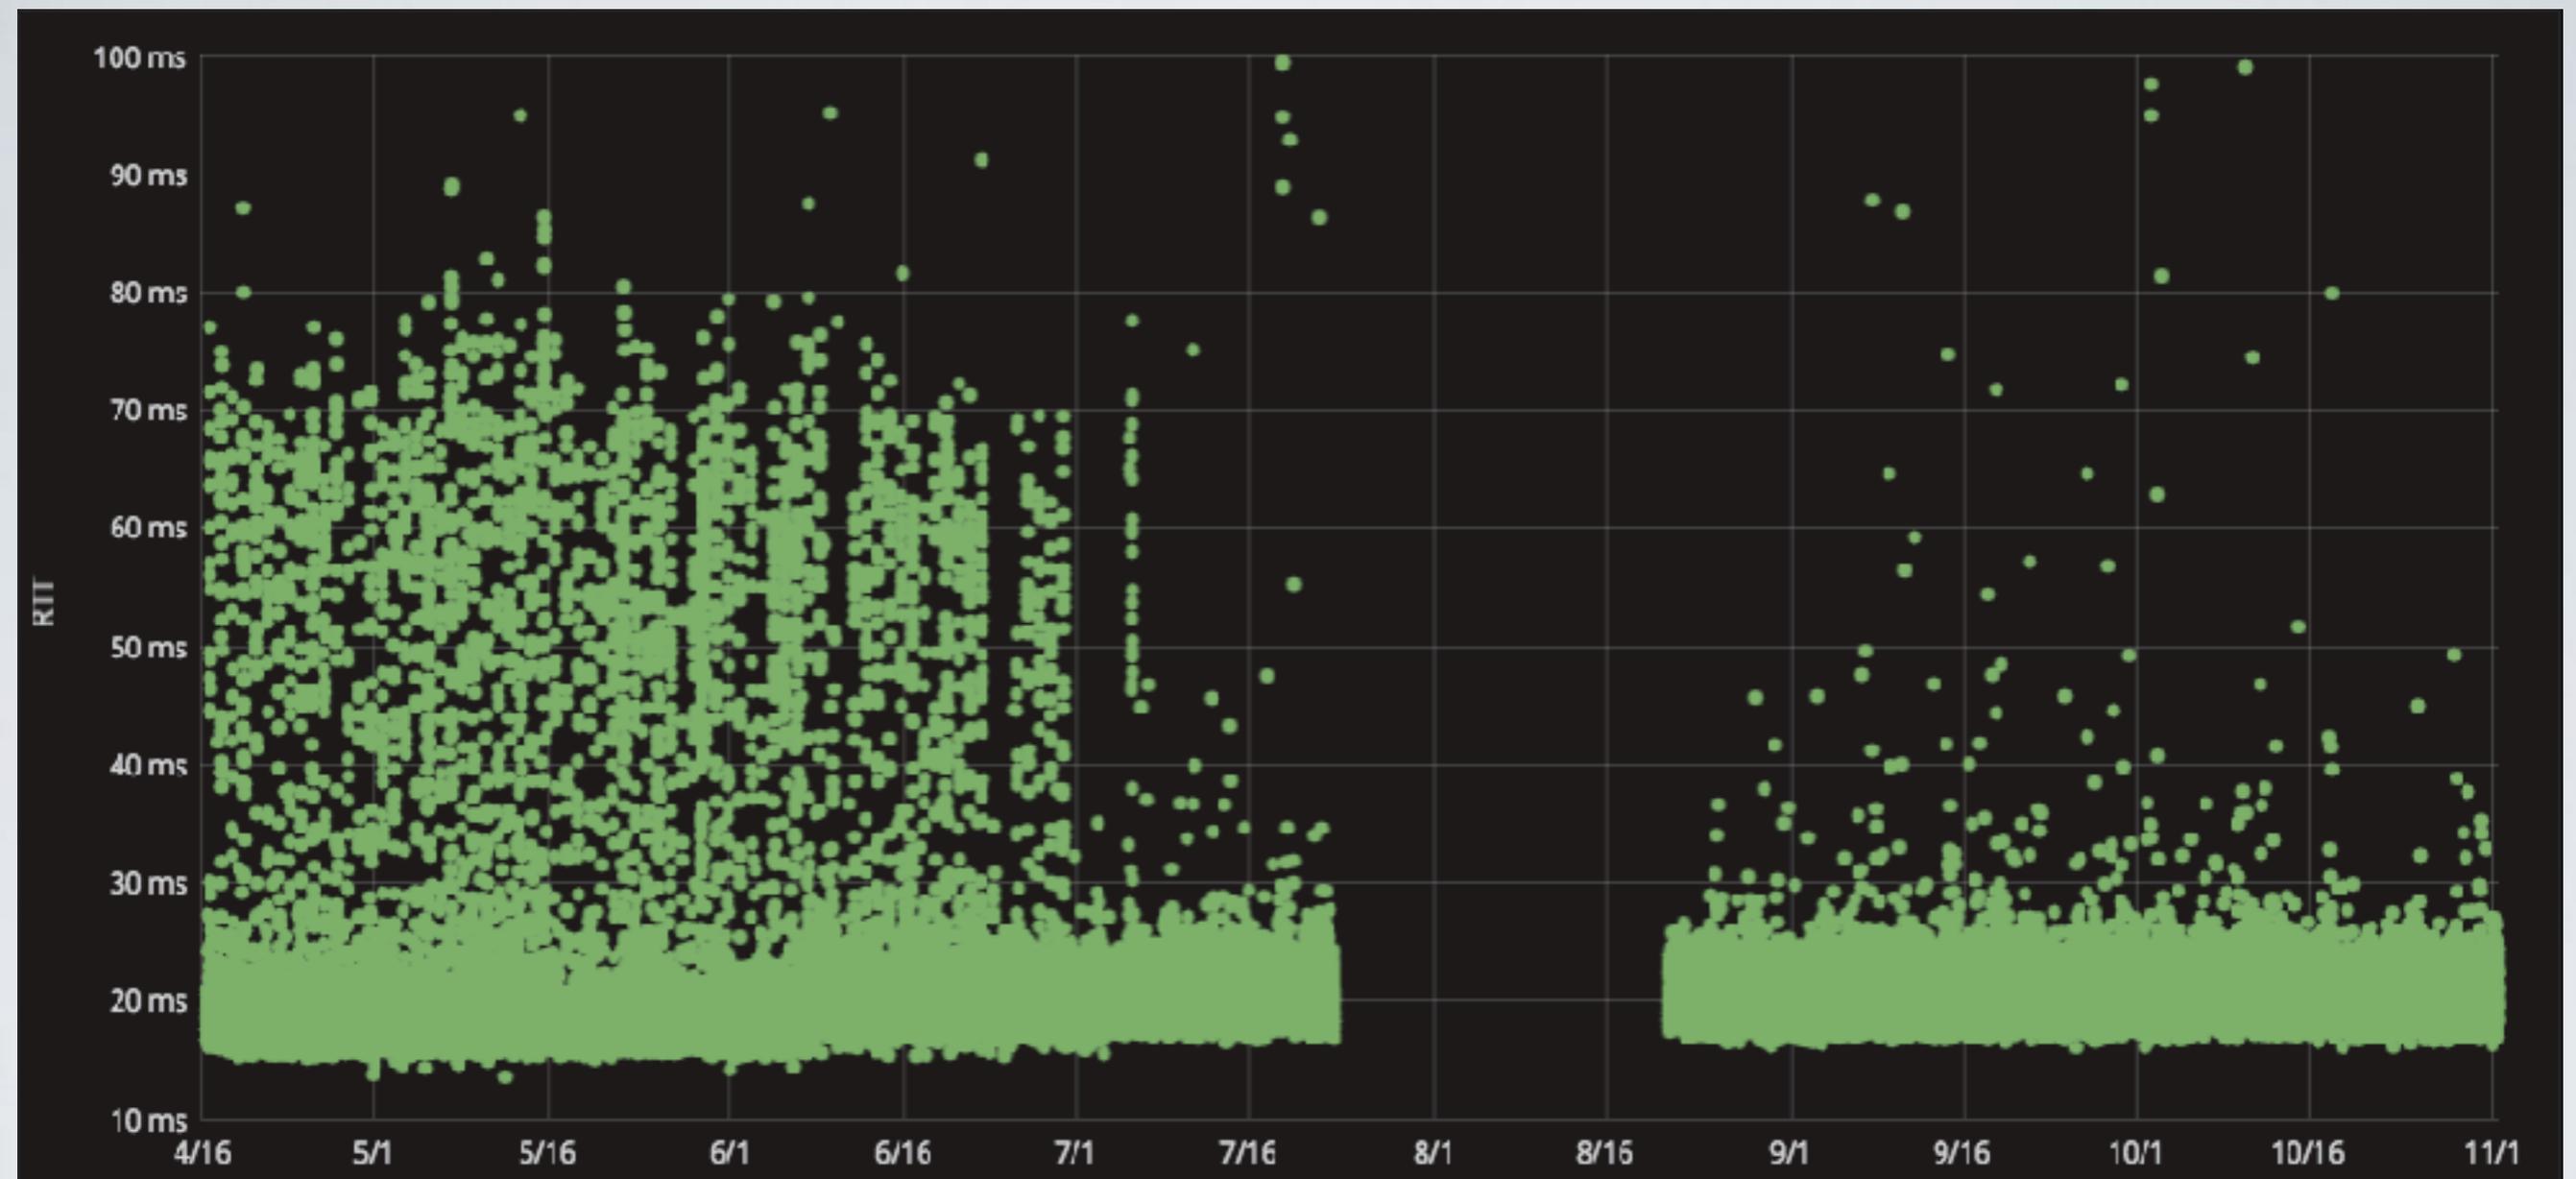
MANIC PANDA



Near real-time dashboards



Longitudinal Views



Comcast - Google, April - November 2017



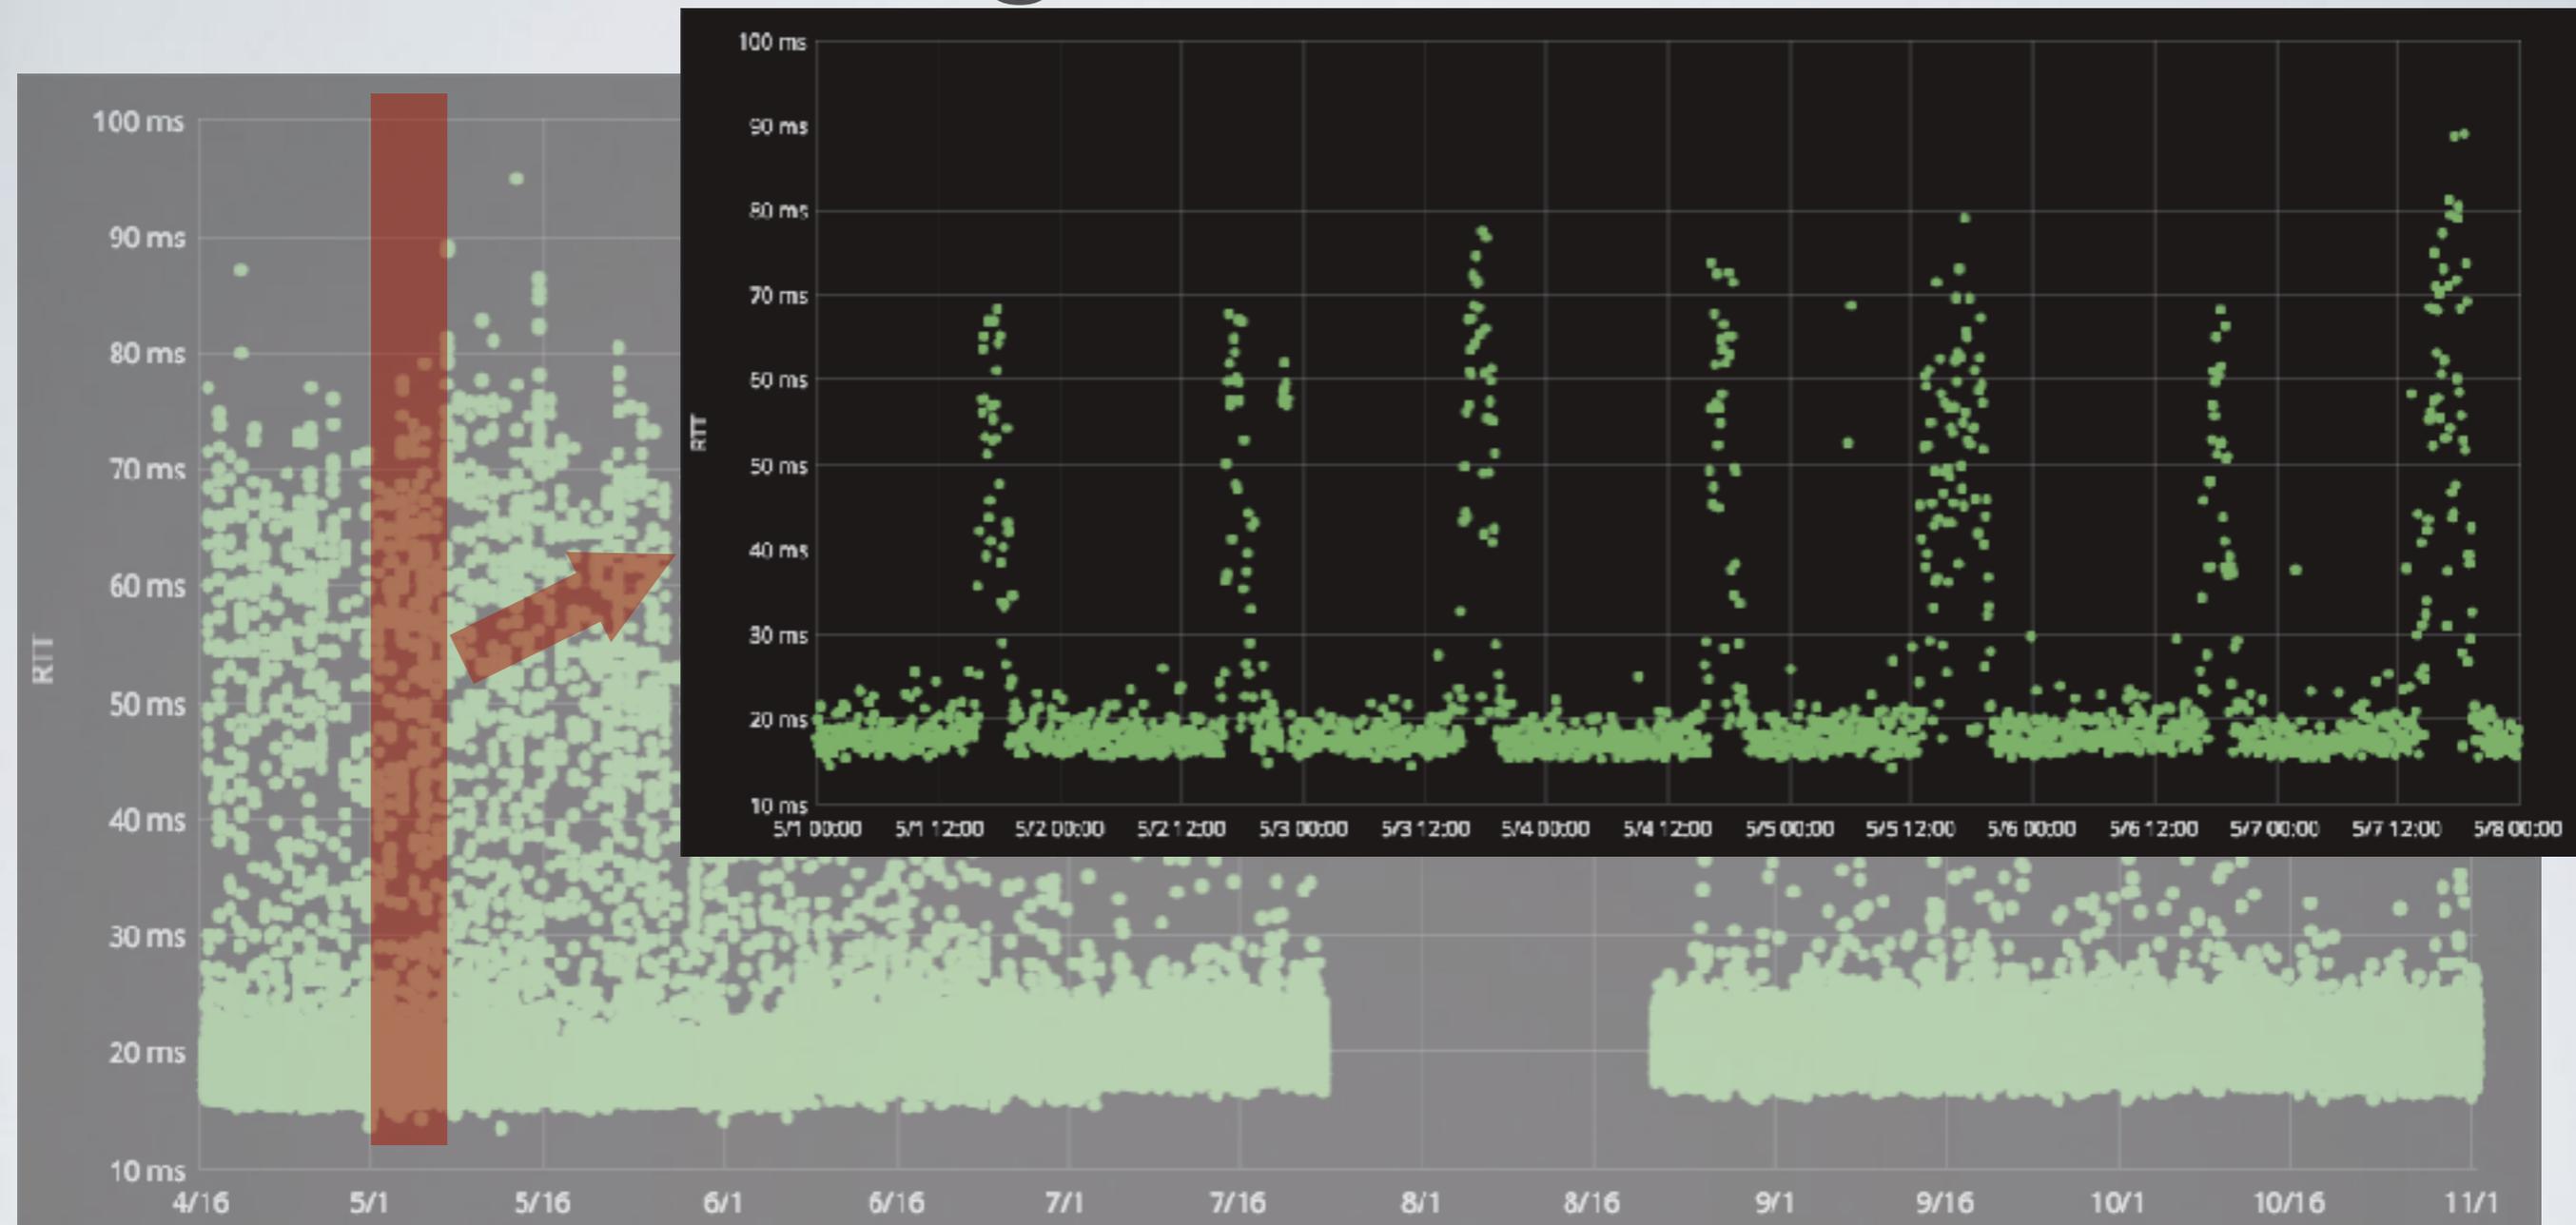
Longitudinal Views



Comcast - Google, April - November 2017



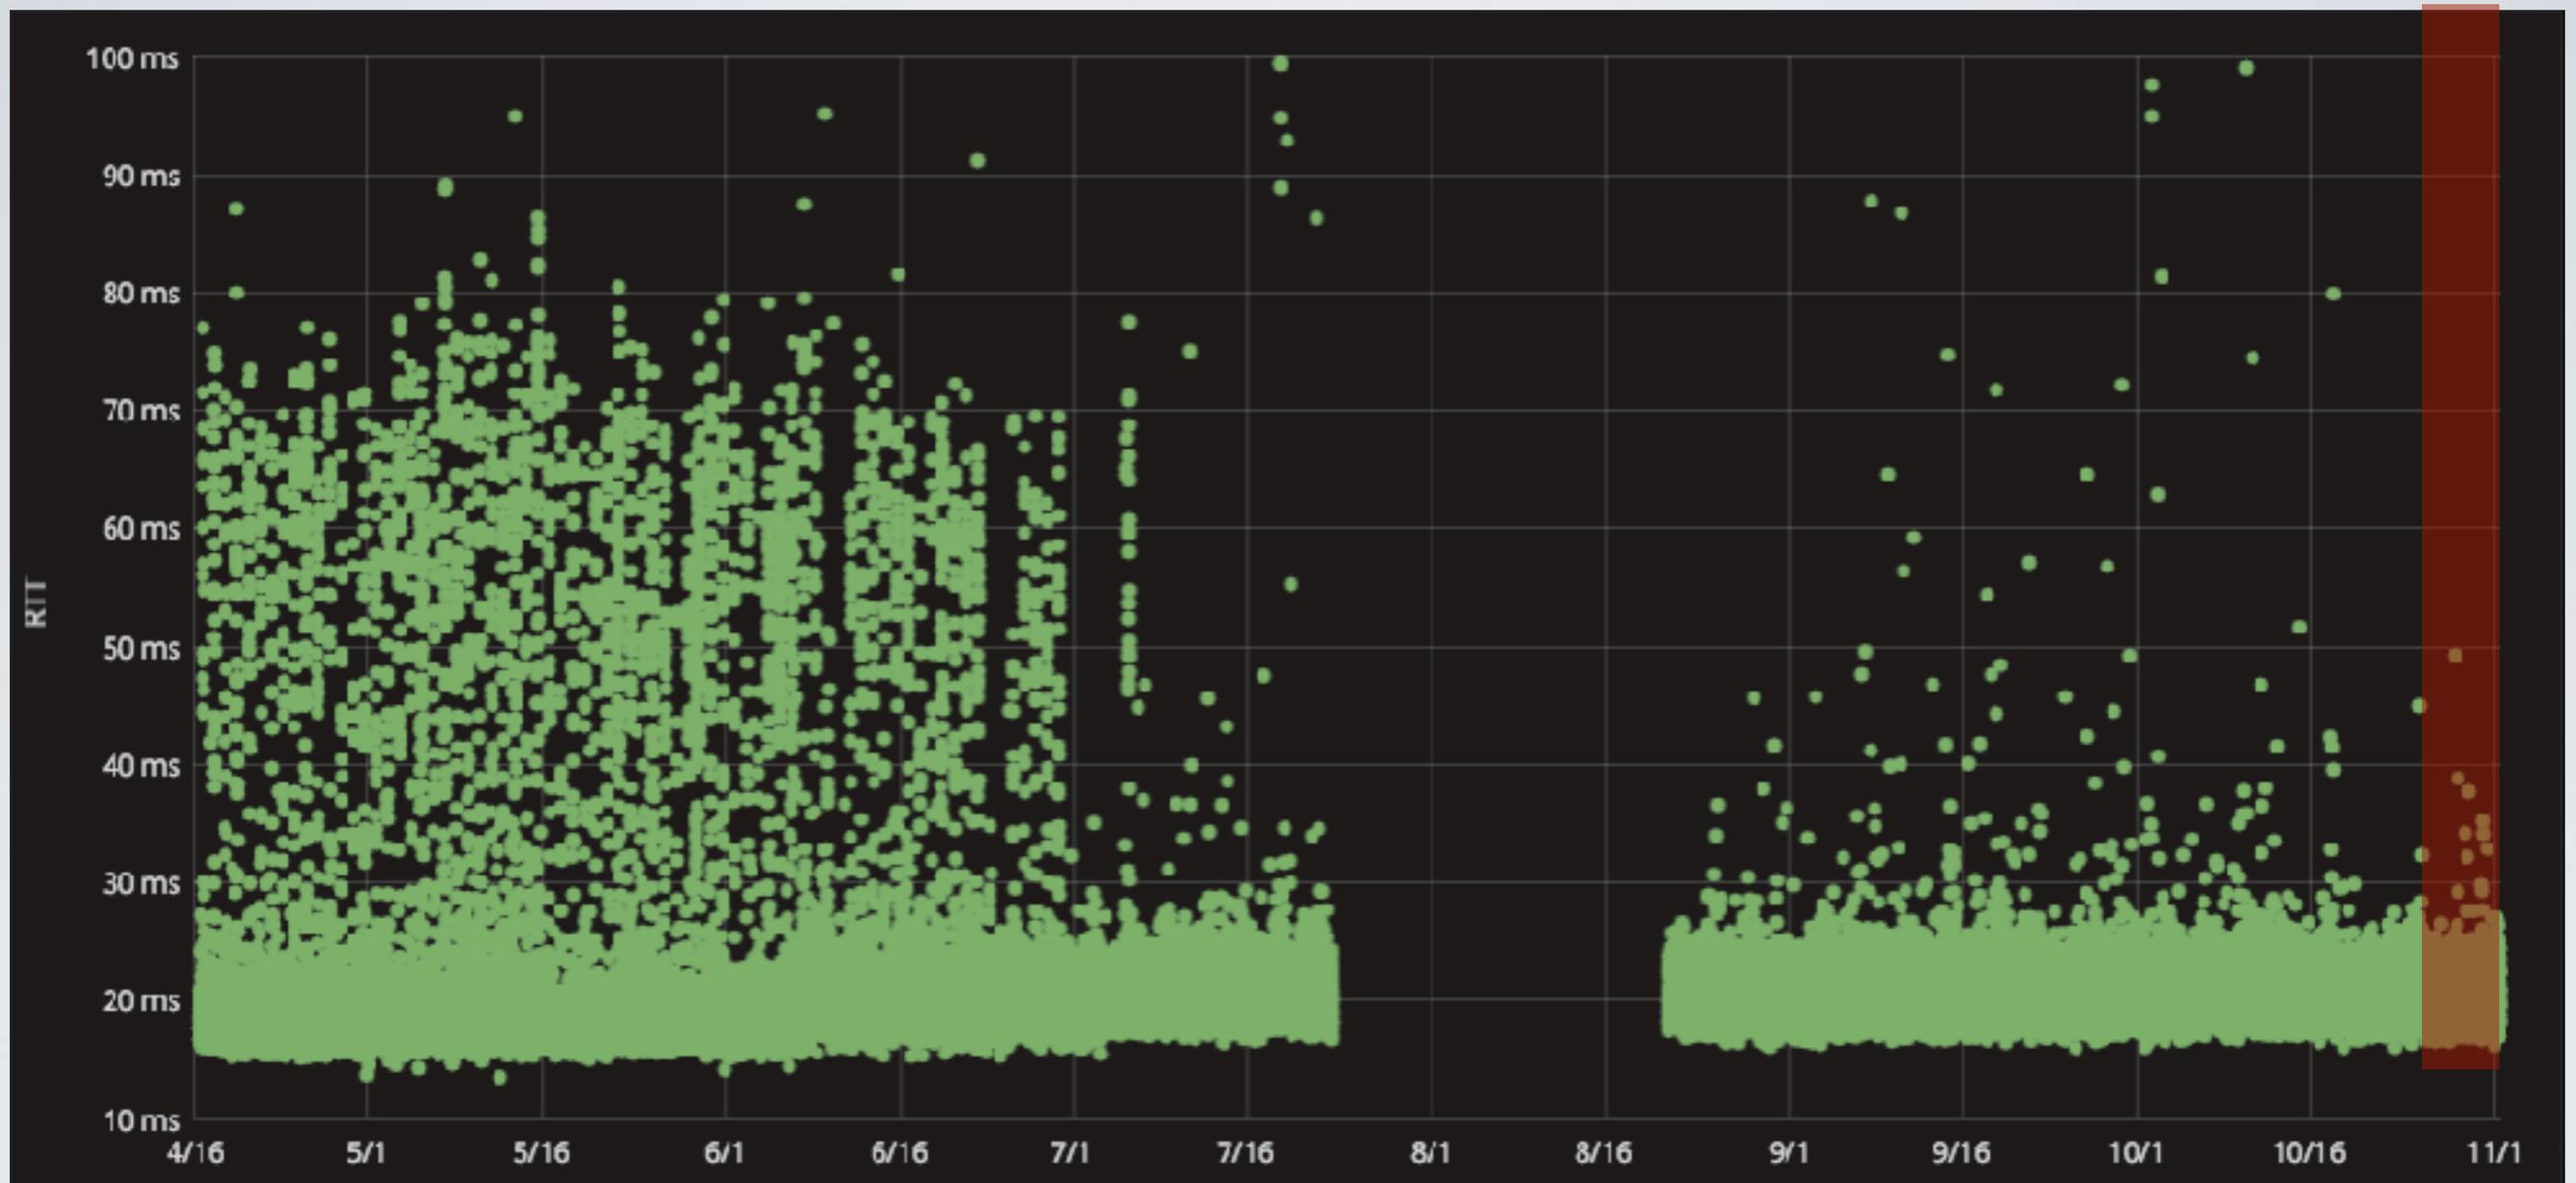
Longitudinal Views



Comcast - Google, April - November 2017



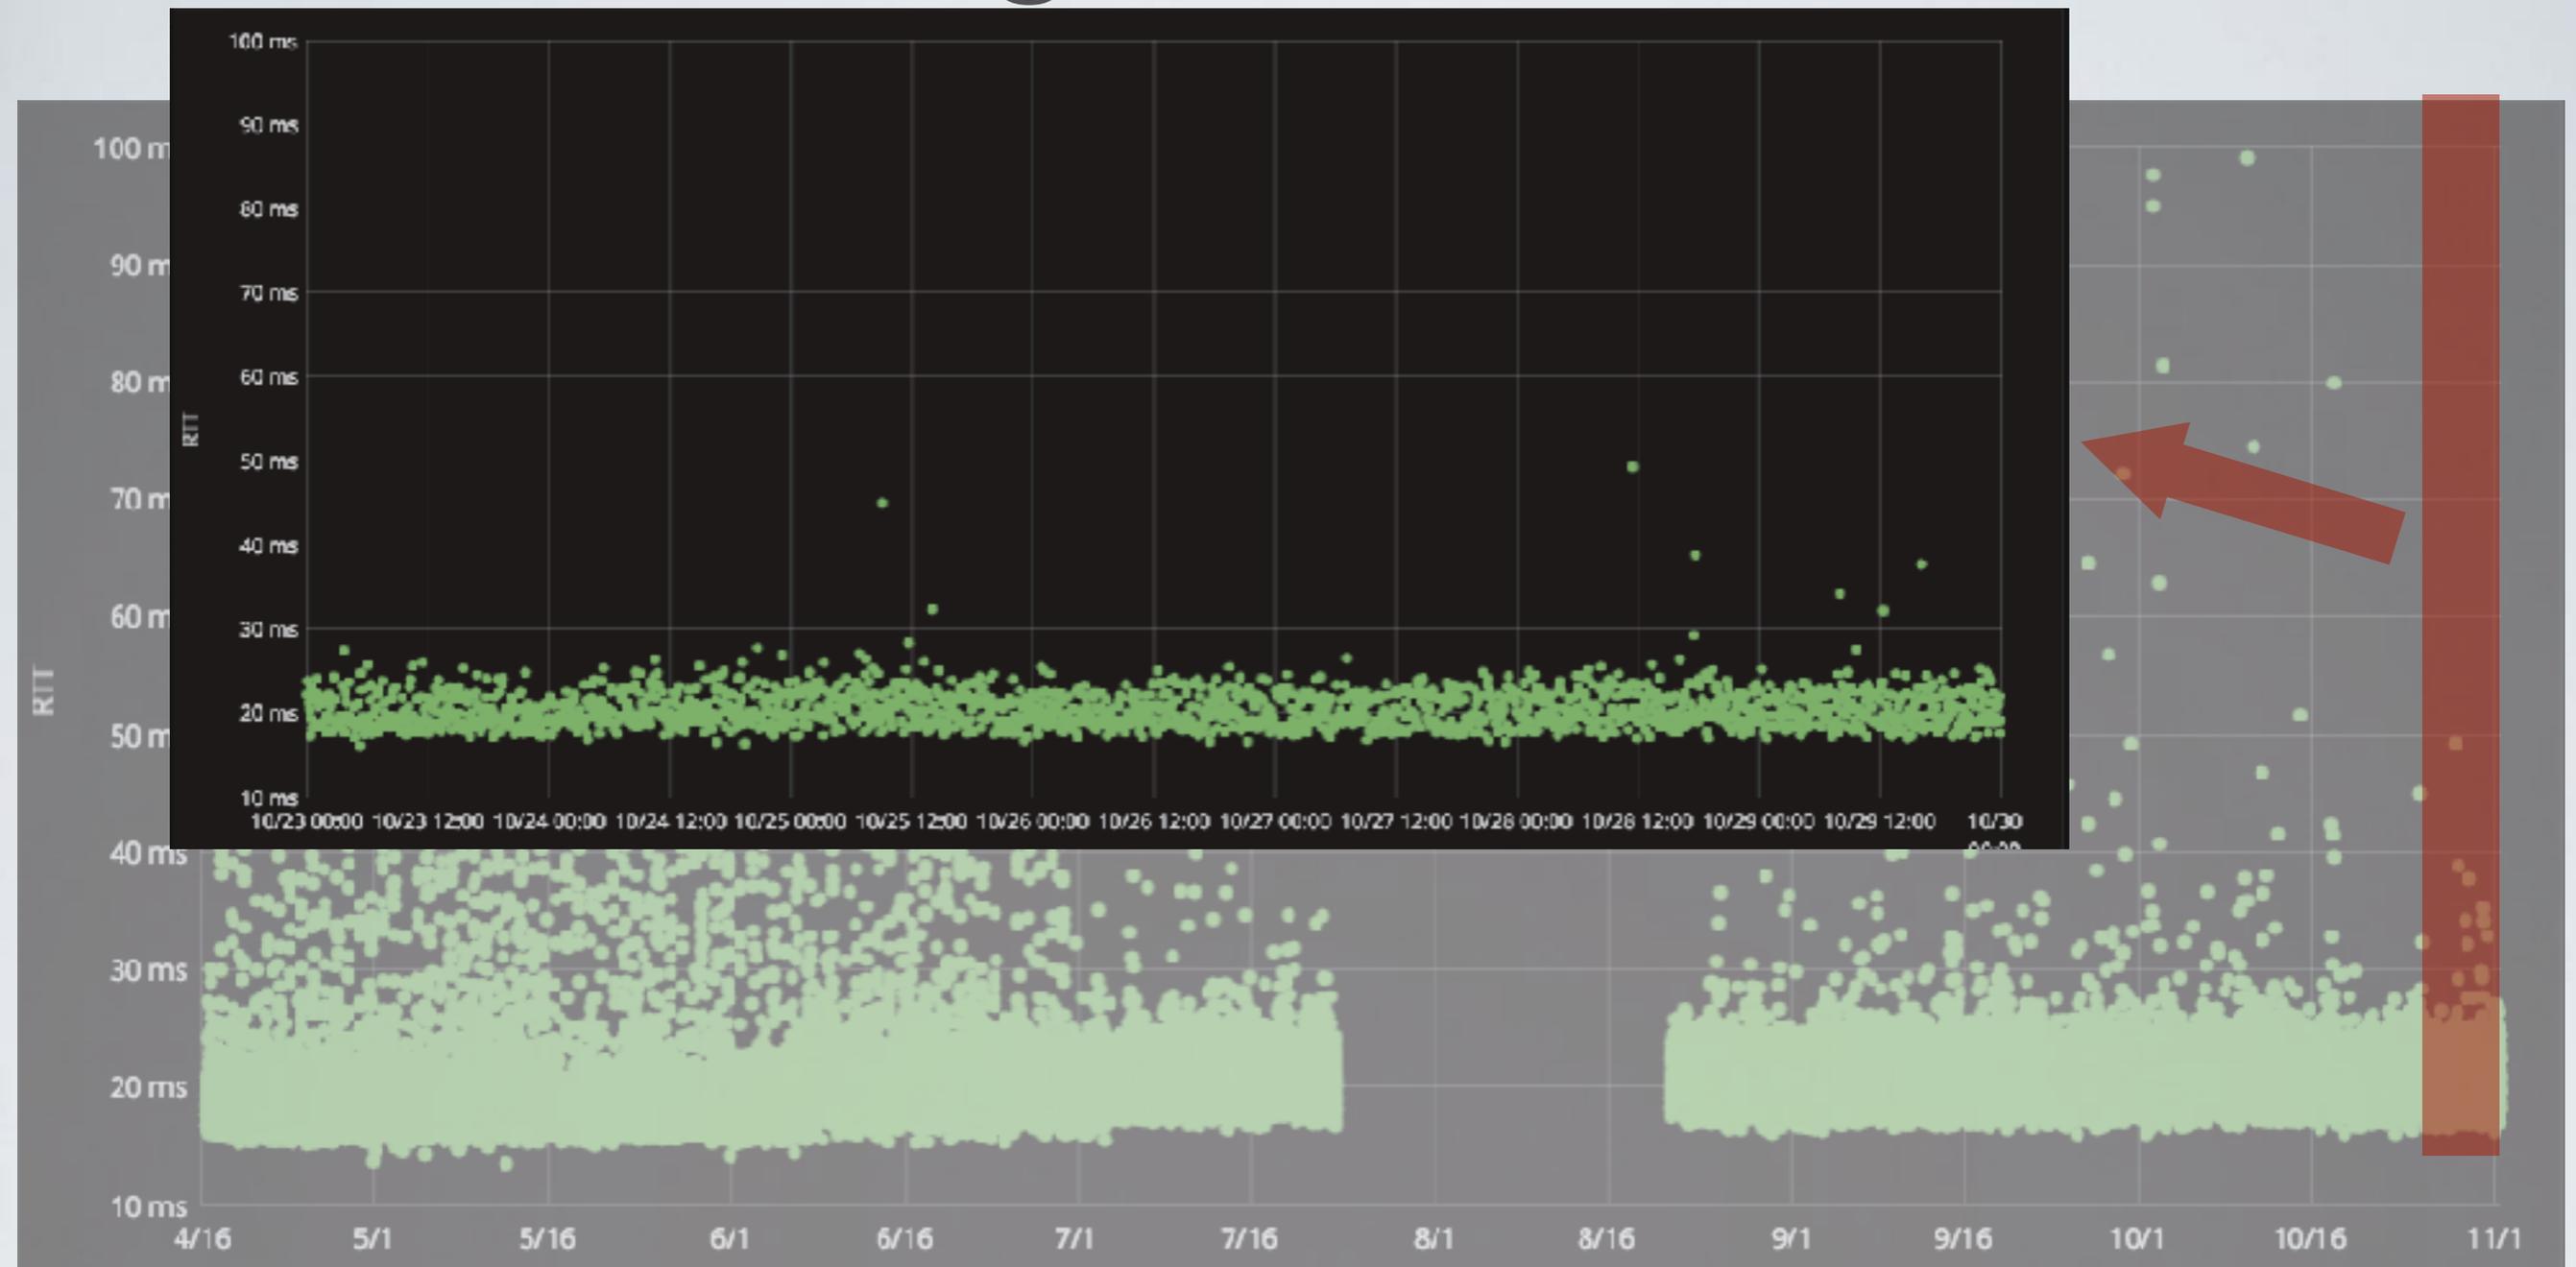
Longitudinal Views



Comcast - Google, April - November 2017



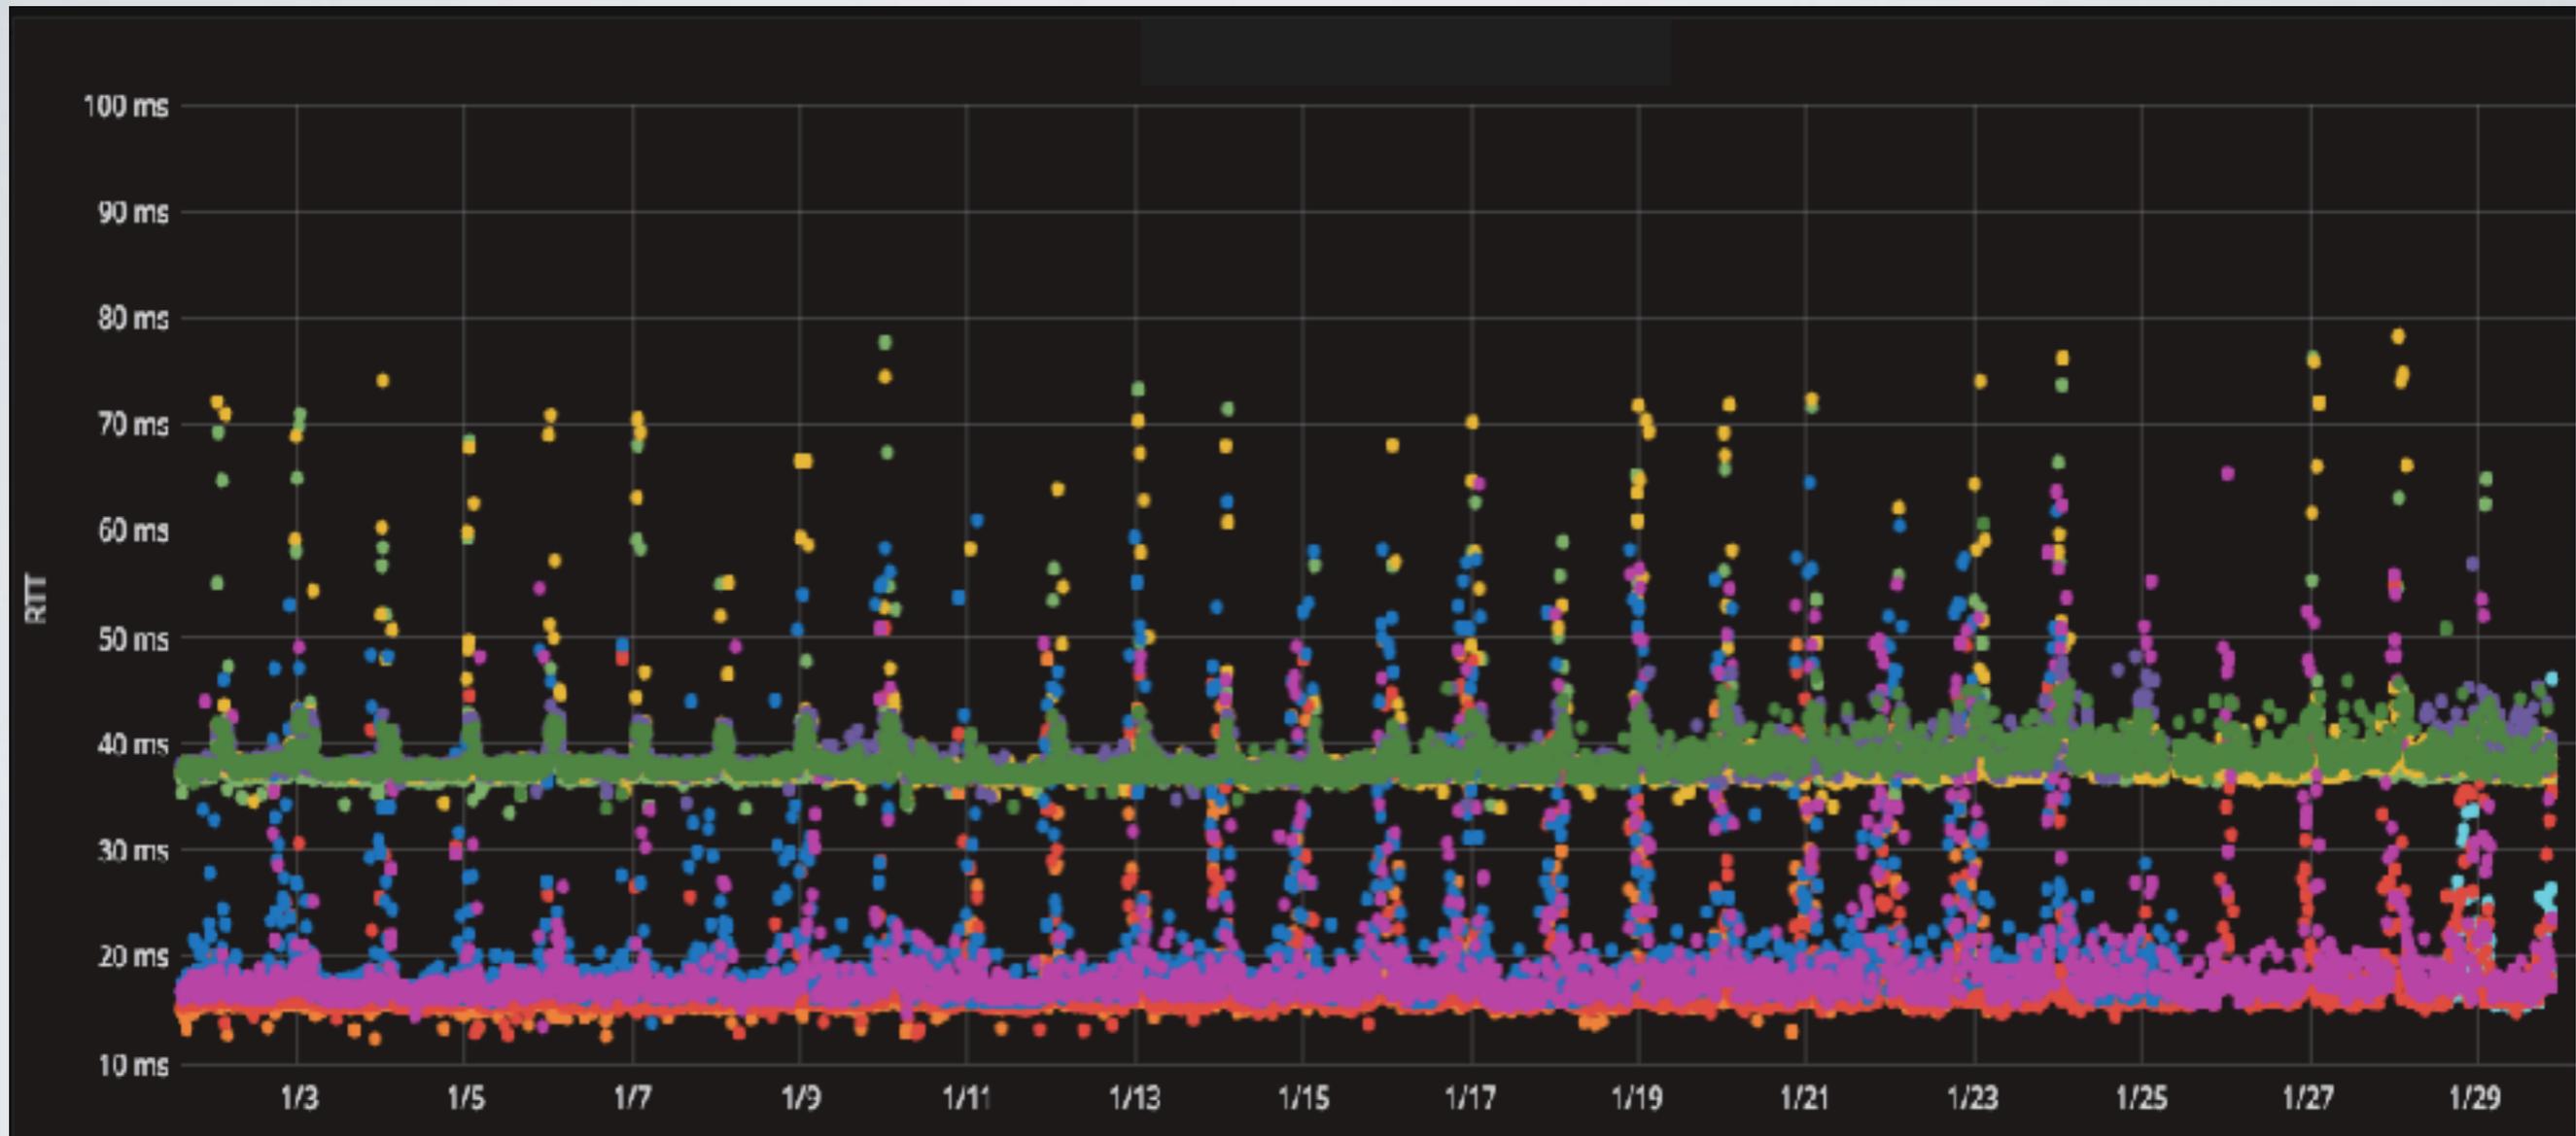
Longitudinal Views



Comcast - Google, April - November 2017



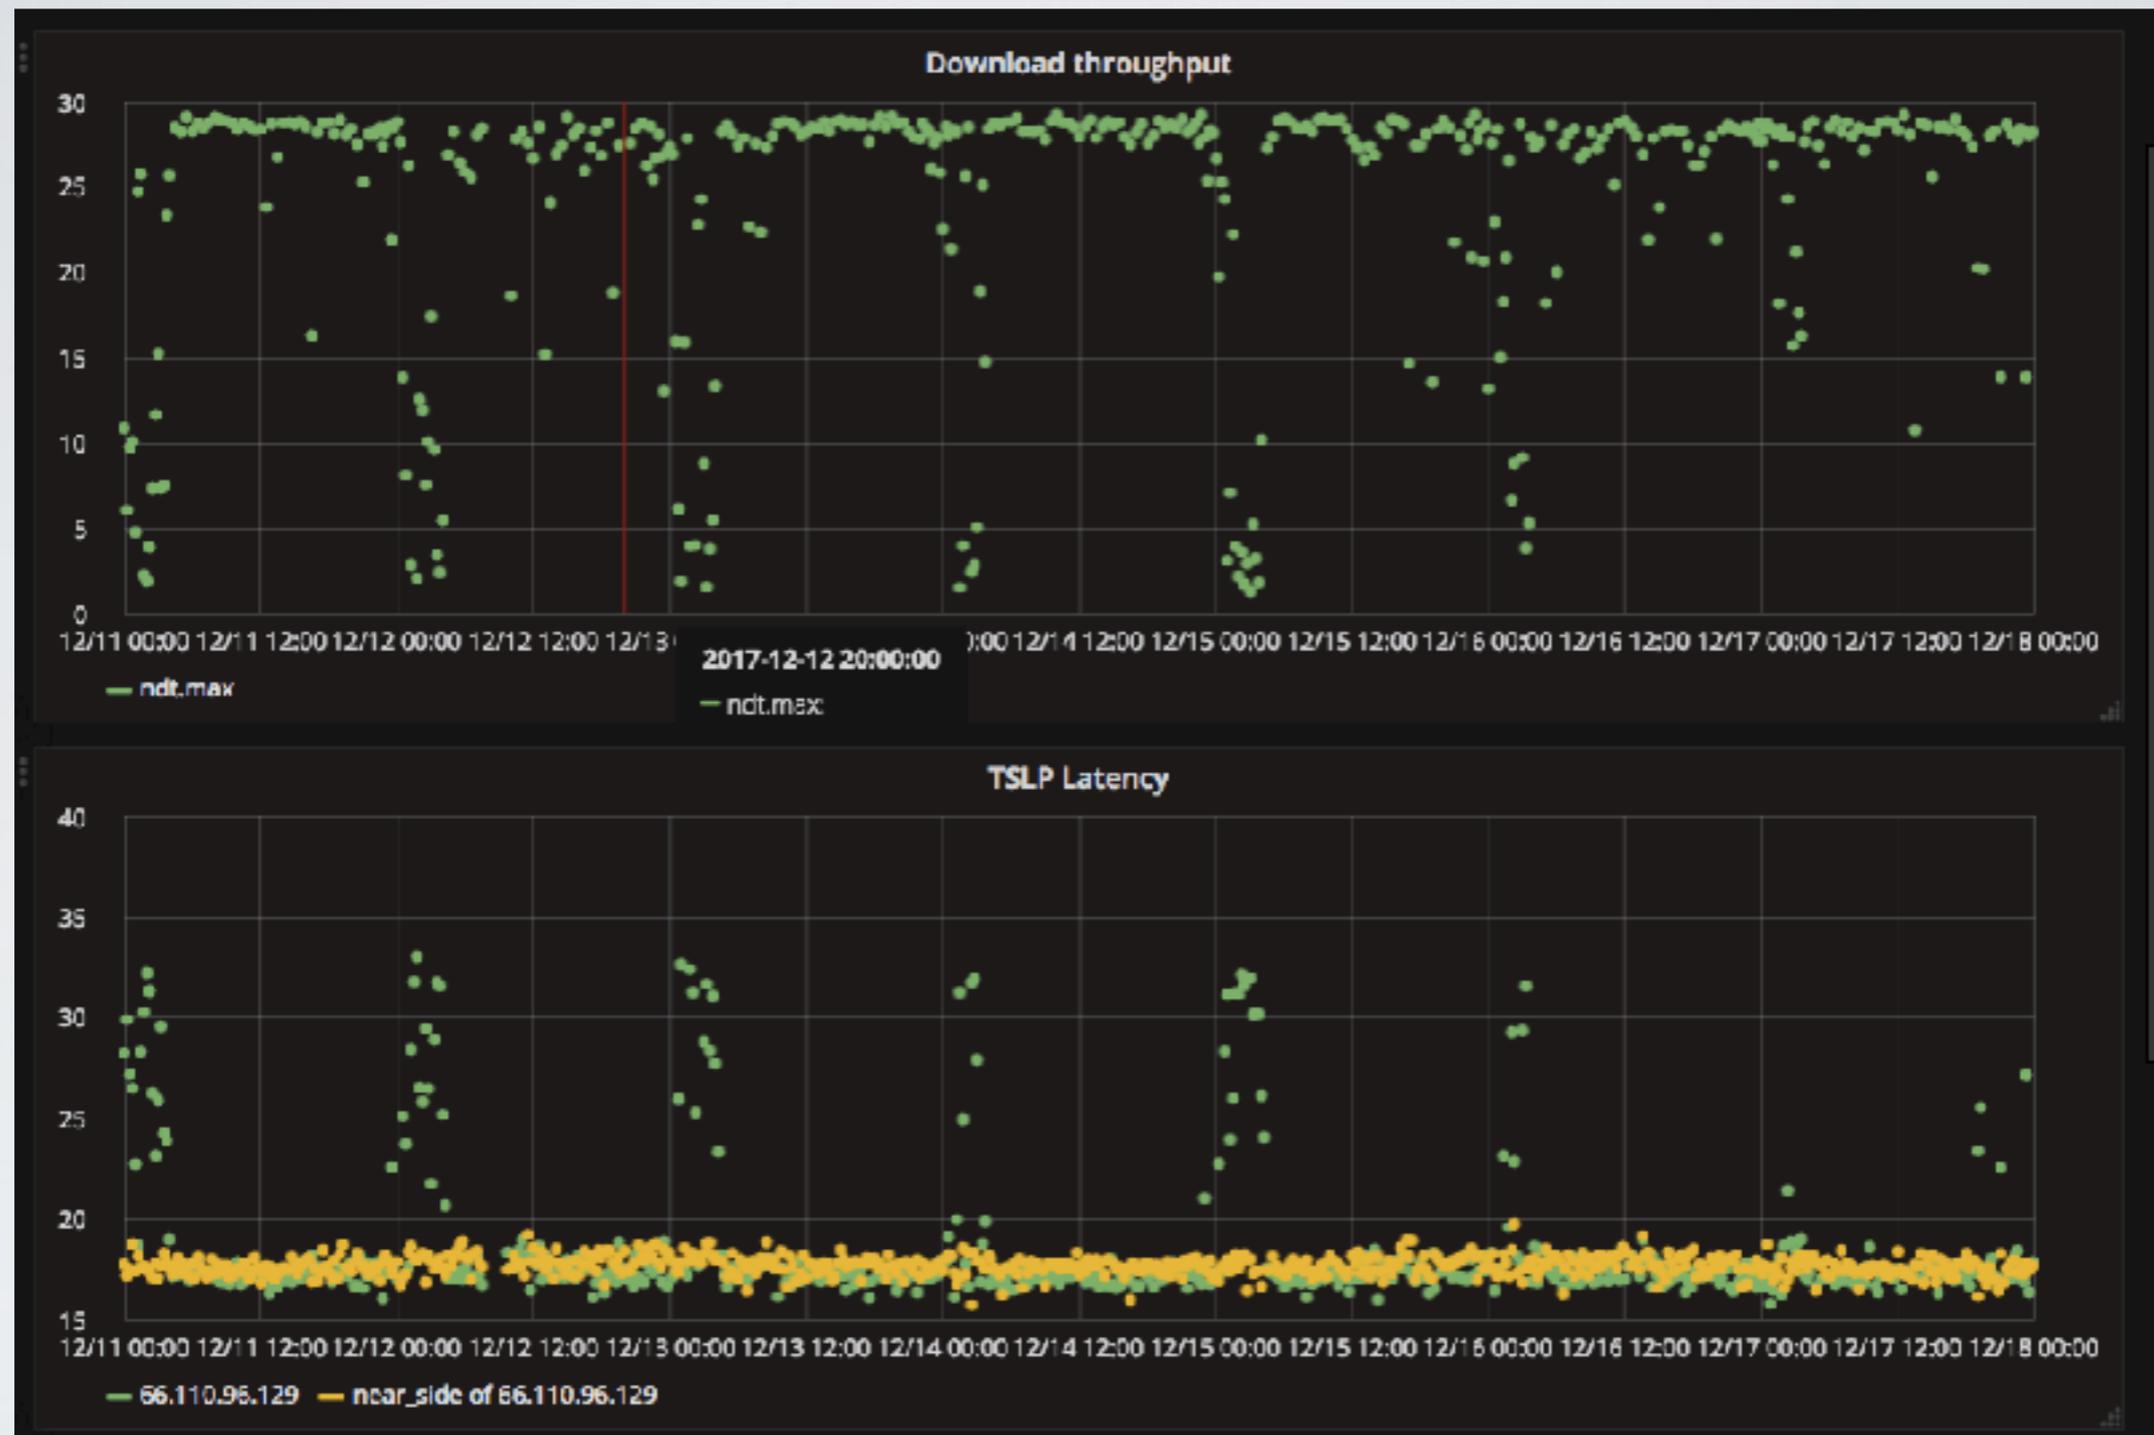
All Visible Links Over Time



Comcast - Google, January 2017



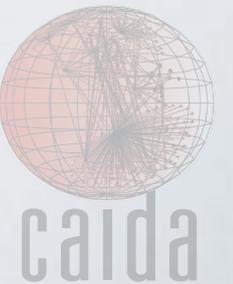
Correlation with Other Datasets



NDT throughput and TSLP latency

Ongoing work in MANIC

- Expanding the coverage of our monitor deployment (always ongoing)
- Improving data quality
- Better techniques for analyzing time-series data, identifying reverse-path issues
- Opening up the web interface for public view, validation from operators
- Investigating impact on other metrics: throughput, loss rate, QoE



Thanks!
Questions?