



Death of the Peering Coordinator

Remco van Mook, November 2021

**You're negotiating with a computer
- how does that make you feel?**



Peering Evolution

The disappearing craft of the peering coordinator

Large scale consolidation

Job openings either at new companies or at strategic levels

The number of people involved hasn't grown with traffic

Interconnection decisions at a higher level

“Hey we're both on exchange X, let's peer” turning into
“please give us space and power across your footprint for our cache node”

Google hasn't had a 'peering coordinator' position for about 5 years



The Tragedy of Least Cost

Interconnection as a cost centre

Only the largest access networks make money on interconnection

The SLA of your transit provider mostly covers availability & congestion

Most are solving for traffic volume, much less for performance



Real Time is Here

..And COVID kicked it forward

Traffic Volume is mostly a solved problem - and usually cheap

Online collaboration is NOT about traffic volume

Caching is not a valid approach

Consumers are now Producers, upstream quality suddenly important

Traffic moving from 'observable' protocols to low-latency, encrypted protocols

50 ms round-trip criteria are commonplace, sub-20 ms is coming



The Internet of Services

Cloud, Serverless & Backend Latency

Identifying what interconnection is underperforming is getting harder

Diversity of interconnection locations is increasing

End user experience only partly determined by the initial traffic destination

A shocking amount of application backends live in AWS US East 1!



How Much Do You Know..

..about your network?

- Your Network Monitoring
- Flow Analysis
- Interface Statistics
- Top talkers (in volume)



How Much Do You Know..

..about destination networks?

- Which direction you're sending traffic for them
- From where you're receiving traffic from them
- Beyond that, not a whole lot



What Goals Do You Have?

Shift from a *volume* based approach to a *value* based approach

“Having great Netflix is nice, but if Teams breaks I have a problem”

“My Social Media team blows up if Roblox doesn’t work”

“My enterprise customers will walk away Oracle is slow”



What Goals Do You Have?

But what does the *business* want the network to deliver?, i.e.

“I want direct adjacencies for 80% of my traffic AND 80% of key applications”

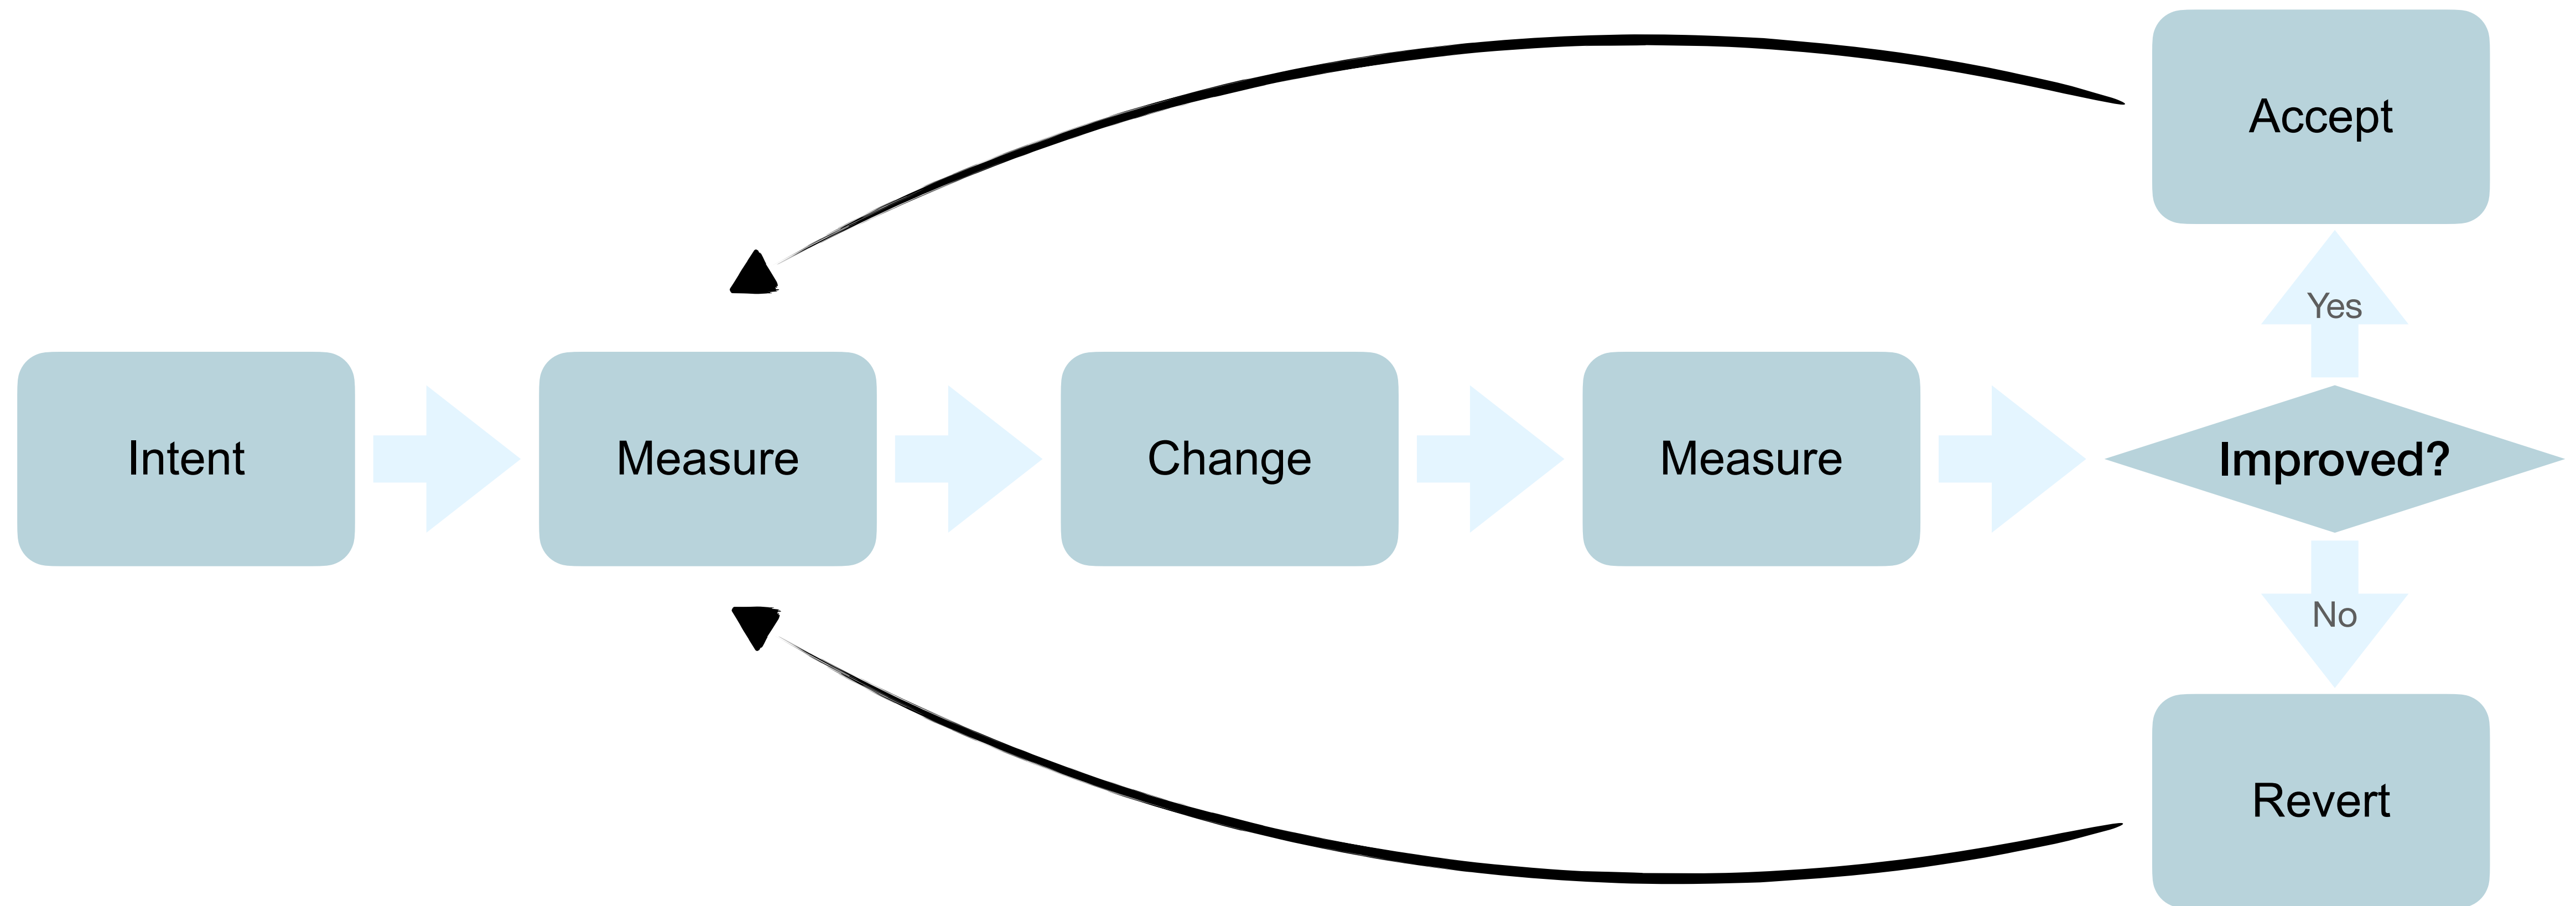
“The top 10 of Cloud/IaaS providers need to be reachable under 50ms”

“End to end packet loss needs to be under 0.01%”



Feedback Loop

How Do You Know You're Achieving Them?



The Trifecta of Performance

1. Latency & Jitter

Round trip time, Variance & Black holes

2. Path

BGP and Trace

3. DNS

Time to First Answer, Black holes (again), Recursive Servers





Measuring Latency



Cloud Measurements

Linkstate end user measurements from the SEE region

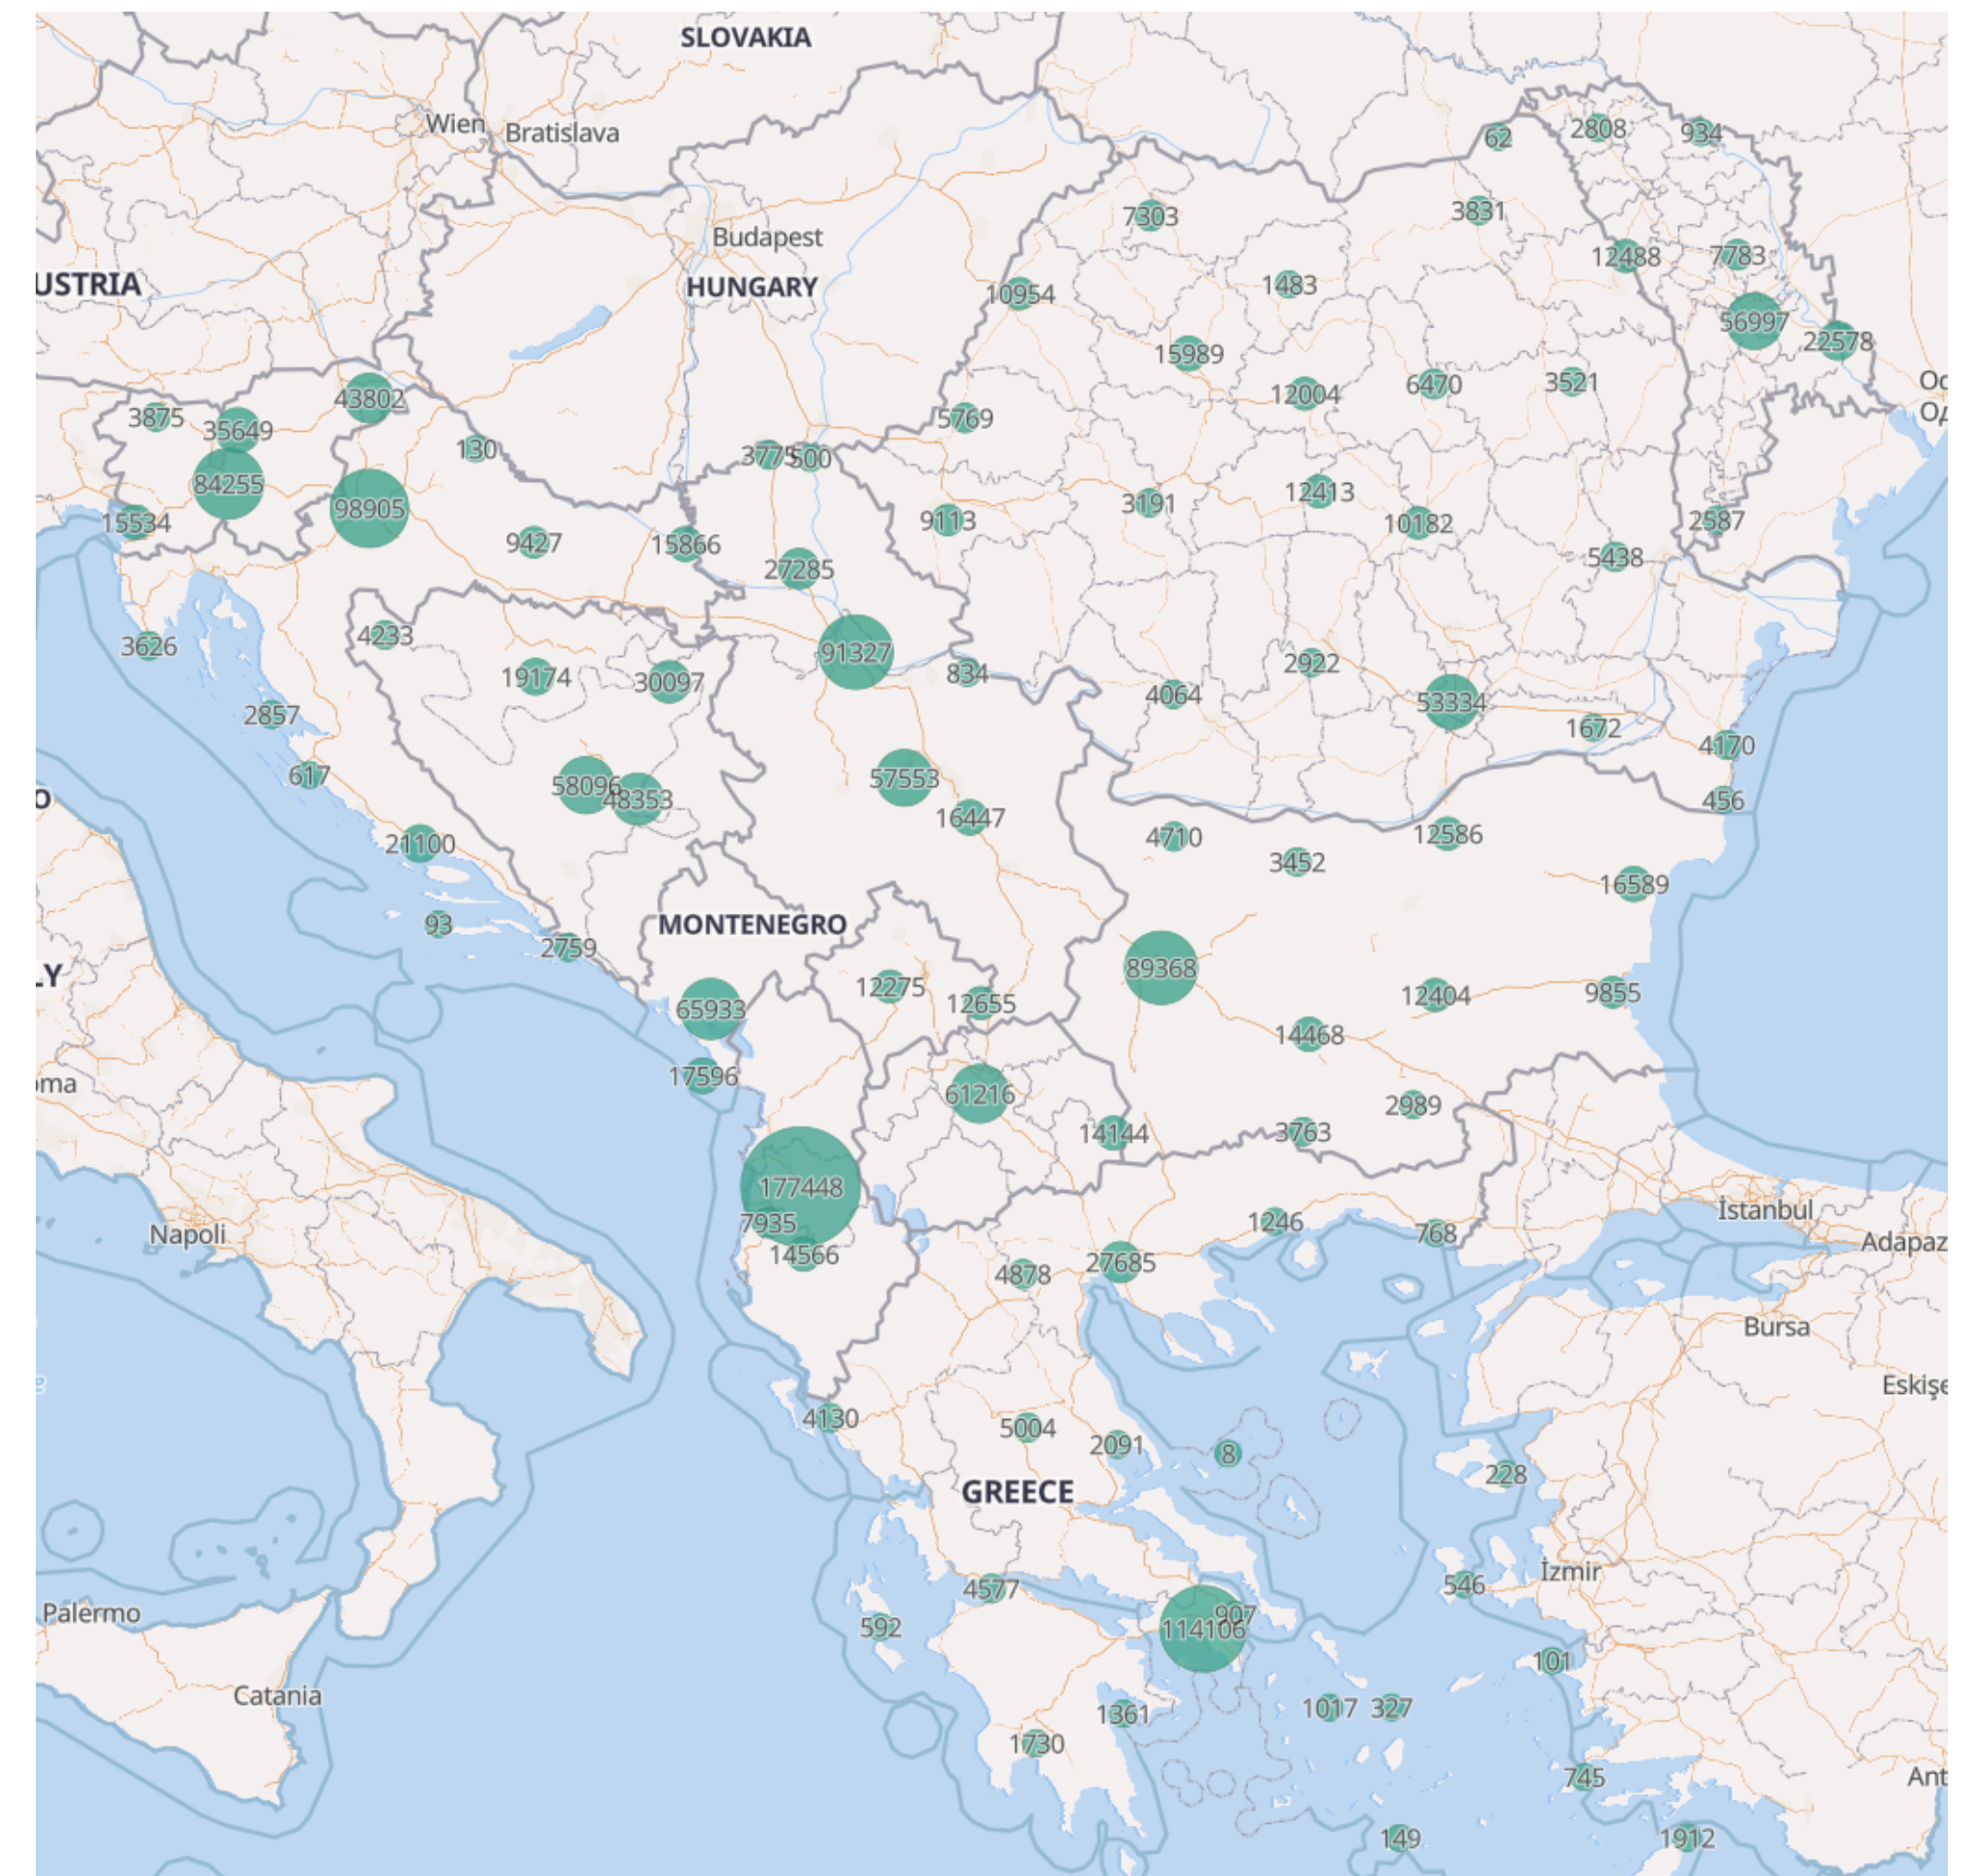
Last 2 weeks

Real end-user connections & devices

~1,800,000 measurements

~120,000 different POVs

(RIPE Atlas has ~400)



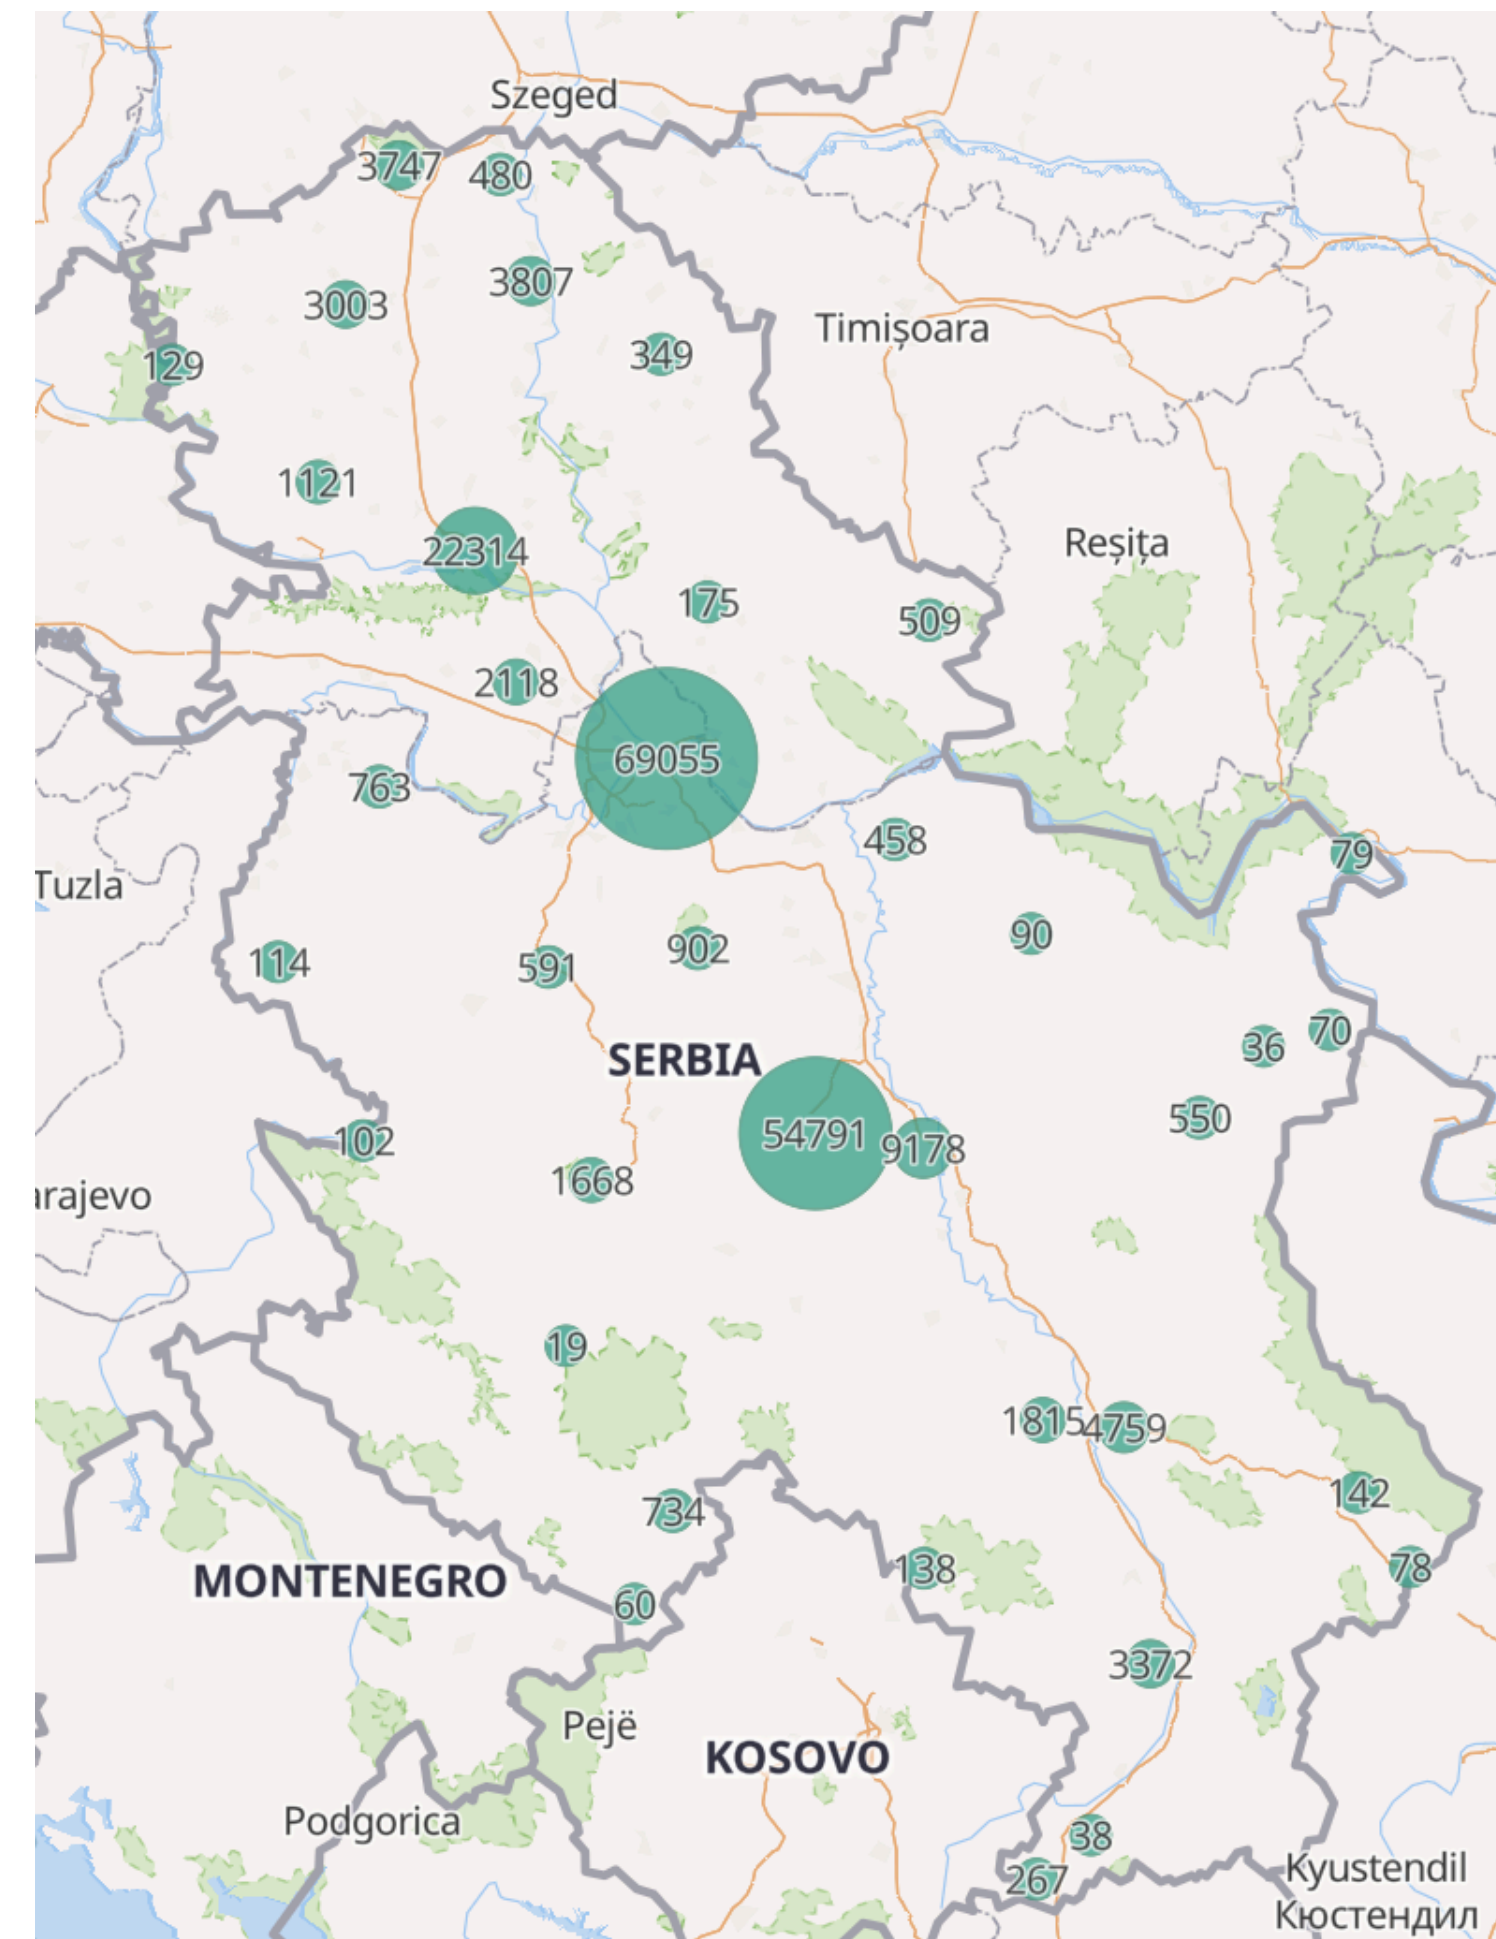
Let's Look at Serbia

Last 2 weeks, all Providers to European Cloud Regions

~200,000 measurements


~12,000 different POVs

~50 different ISPs/ASNs




Let's Look at Serbia

Last 2 weeks, all Providers to European Cloud Regions



Let's Look at Serbia

Last 2 weeks, best performing Cloud Regions




Frankfurt wins despite being further away

Google Cloud is (marginally) the faster provider



Let's Look at Serbia

The Look of Latency



50% of performance is fine


25% of performance is ok

10% of performance is 'ehh'



Someone in Serbia..

Is making weekly backups to AWS Frankfurt every Tuesday





Measuring Path

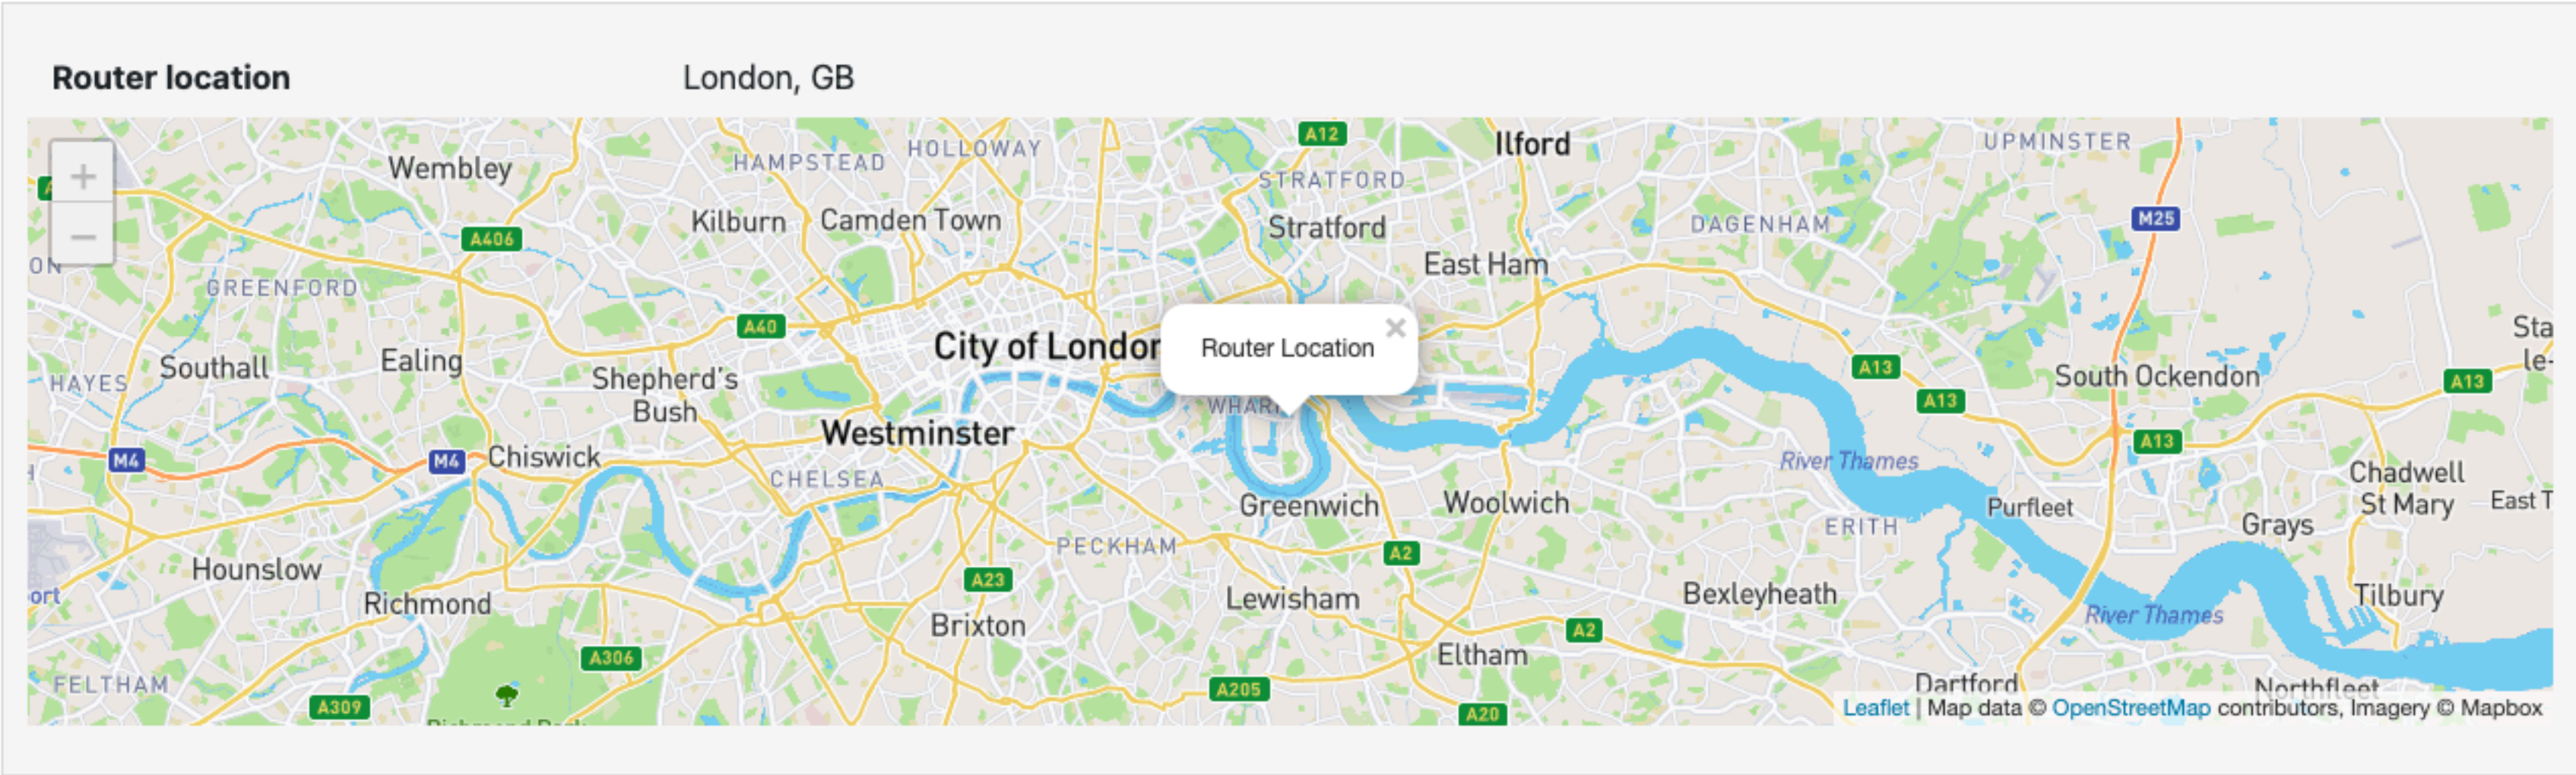


TrackBGP alarm for: AS 51088

Details

Time of Alarm	Sat Nov 13 2021 13:01:50 GMT+0100 (15 days ago)	Resolved
Type of Alarm	invalid_transit	
Origin ASN	60557	
Observed path	44684 1299 174 58299 8298 50869 51088 50673 44103 60557	
Problem	AS50869 is not specified as a transit provider for AS51088, but we saw BGP updates with this transit adjacency.	

Last seen



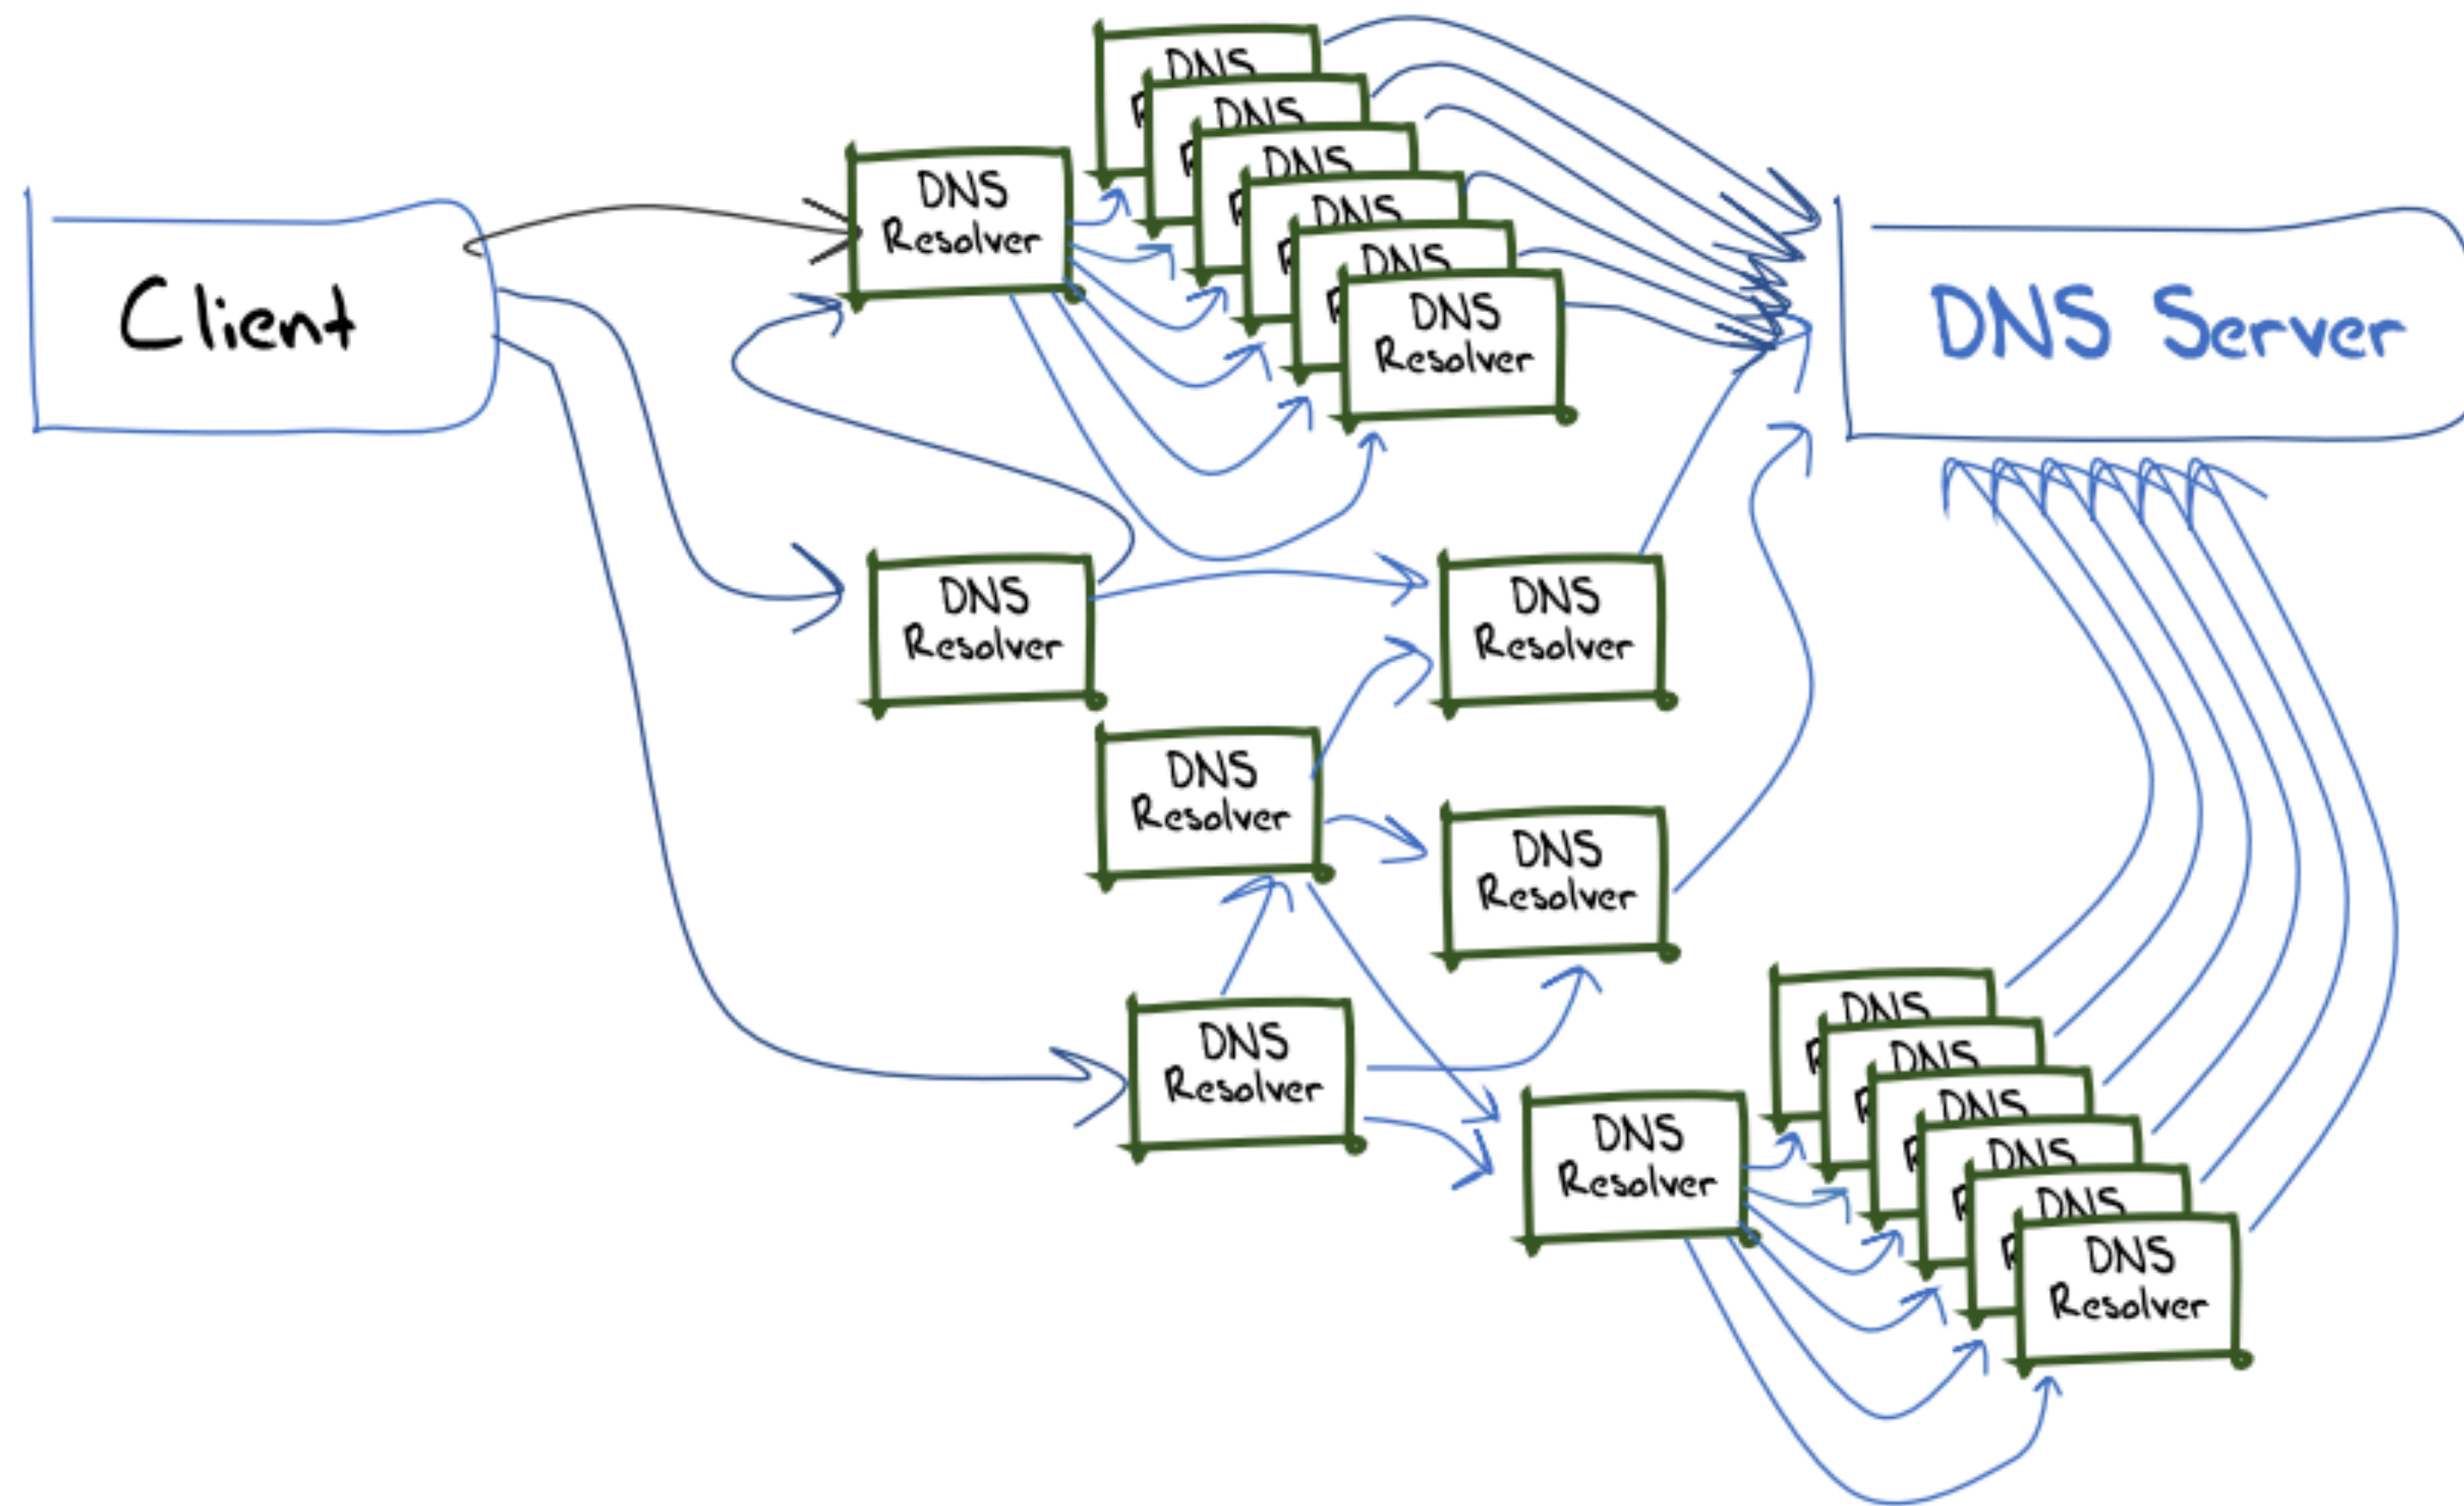


Measuring DNS



The DNS “cloud”

From Geoff Huston’s presentation at RIPE83



Why is this Important?

- ✧ DNS has become the ‘service’ directory
- ✧ DNS outcomes vary based on which DNS resolver gets used
- ✧ The success rate of a *well-configured* DNS setup is only 98%
- ✧ “Wrong” answers can nullify any efforts made in network optimisation
- ✧ DNS caching is a double-edged sword
- ✧ End users with external DNS providers are a pain to optimise (even for CDNs)





Conclusions



Conclusions

Interconnection is no longer just about volume and next hop

Applications moving to cloud settings require a different approach

Automation is the only scalable way forward

The things you're currently measuring aren't telling the whole story

You need high fidelity, large scale network measurements to automate





Questions?

remco@lynkstate.com

www.lynkstate.com

www.trackbgp.com