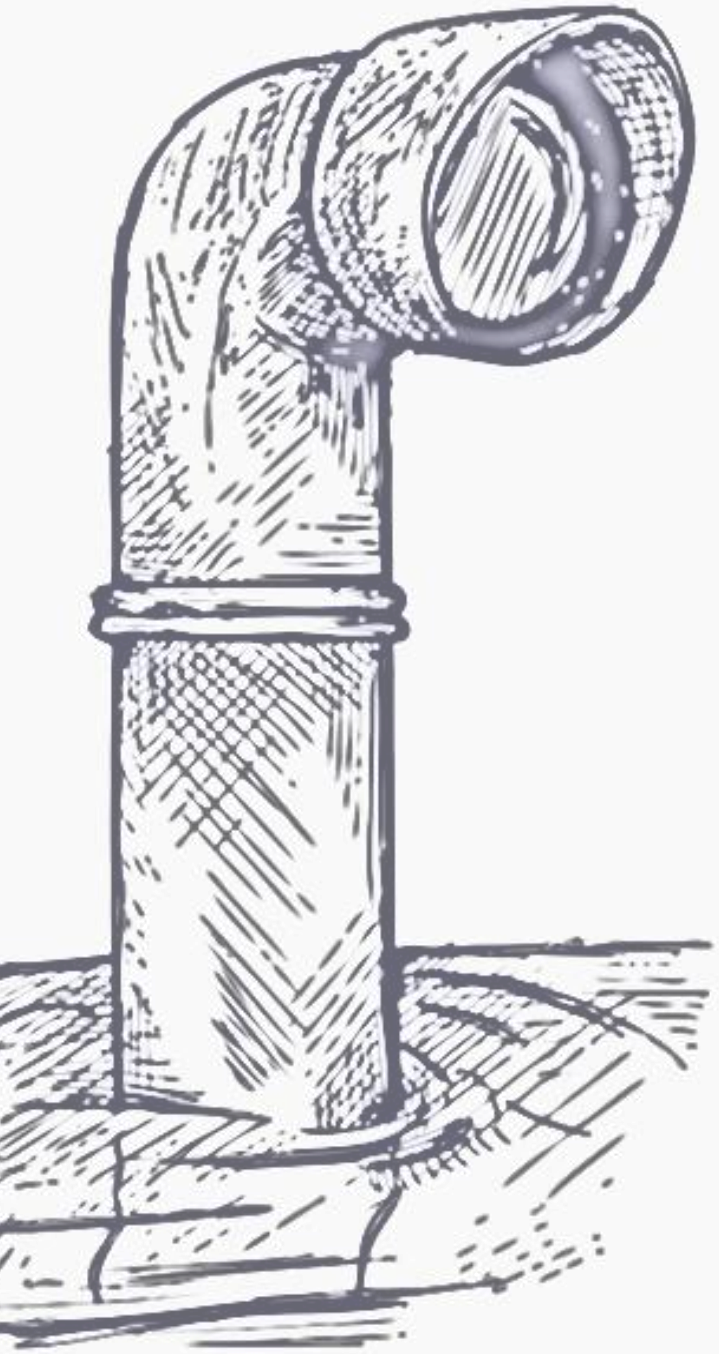

Periscope: Unifying Looking Glass Querying

Vasileios Giotsas <vgiotsas@caida.org>

Amogh Dhamdhere

kc claffy

UCSD/CAIDA

Passive and Active Measurement Conference (PAM) 2016, Heraklion, Crete, 31 March -1 April, 2016

What is a Looking Glass (LG)?

- Web-based interface that allows the execution of read-only network commands
- Provide access to a single or multiple locations within an AS
- Ad-hoc use of LGs in measurement studies

Motamedi, R., Rejaie, R., & Willinger, W. (2015). A Survey of Techniques for Internet Topology Discovery. *Communications Surveys & Tutorials, IEEE*, *17*(2), 1044-1065.

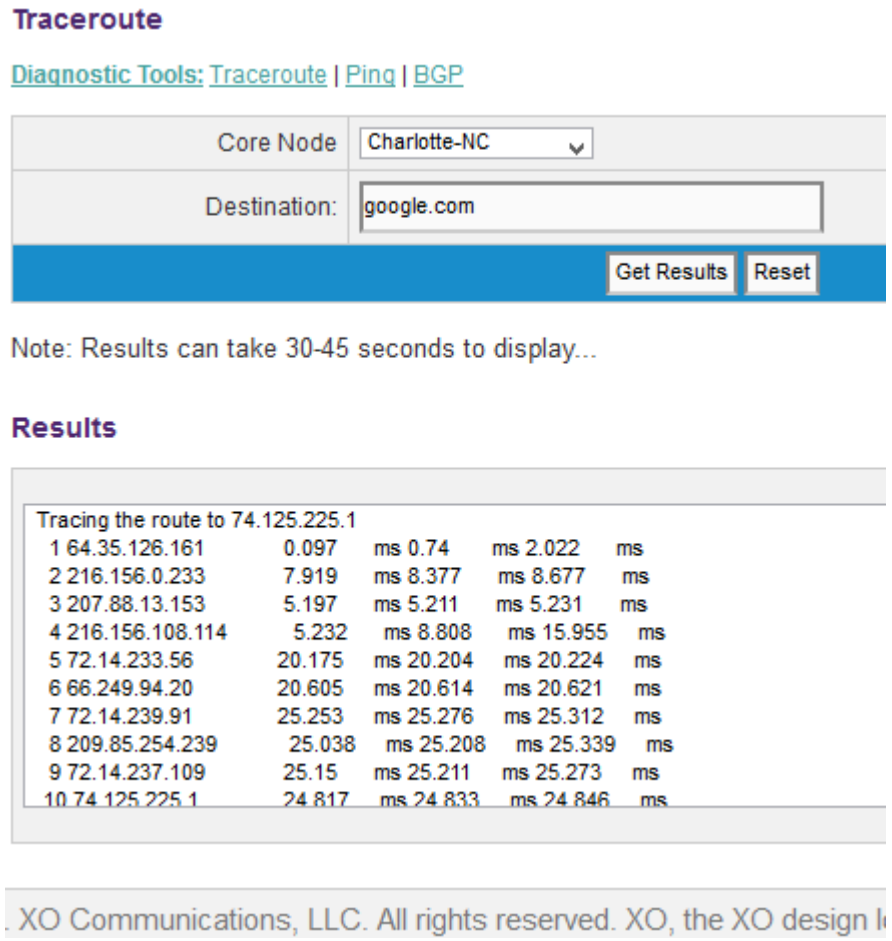

Challenges in the systematic use of LGs

- No authoritative repository of available LGs
- Disparate input interfaces and output formats
- LGs are volatile in terms of availability and specifications
- LGs are intended for low-frequency querying

Requirements for automating LG querying

- No authoritative repository of available LGs
  - Automatically discover LGs and extract of LG specifications
- Disparate input interfaces and output formats
  - Standardize input and output format
- LGs are volatile in terms of availability and specifications
  - Detect status changes and interface updates
- LGs are intended for low-frequency querying
  - Support multiple users without violating rate limitations

Looking Glass discovery workflow

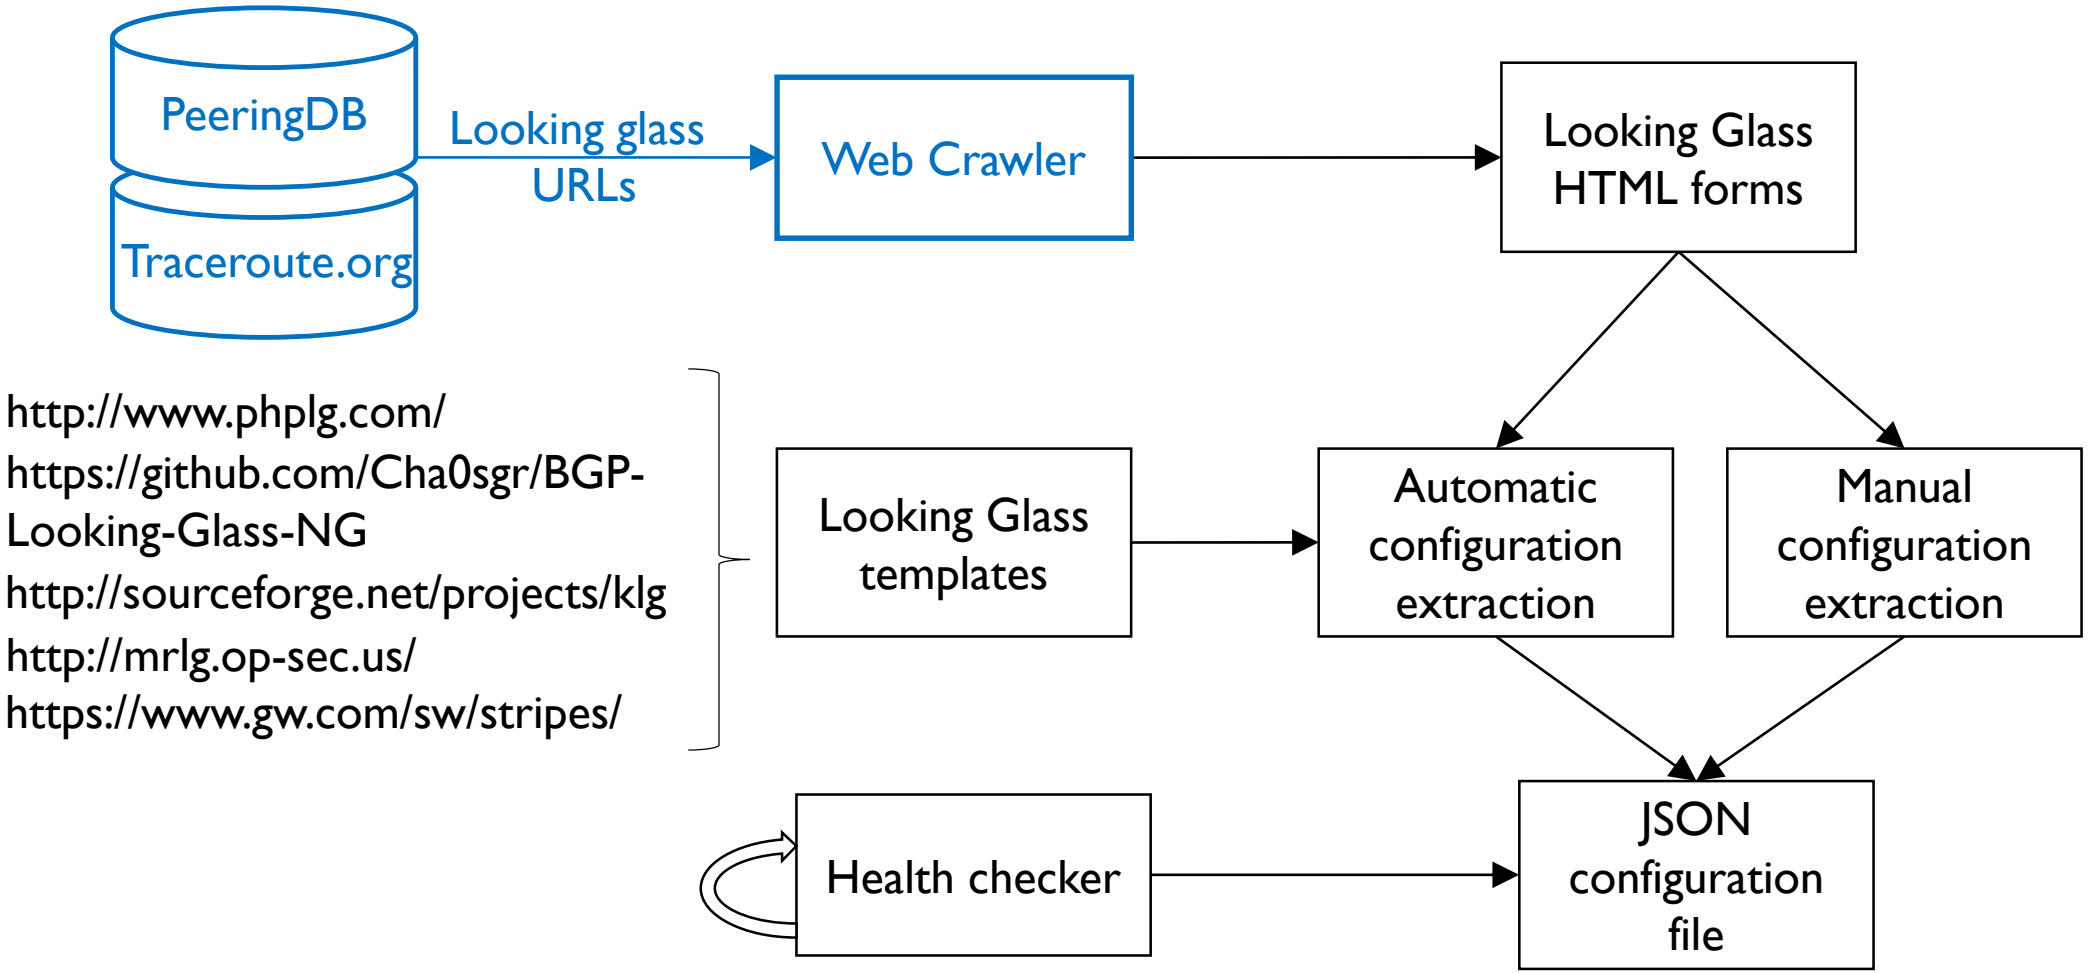

Looking Glass discovery workflow

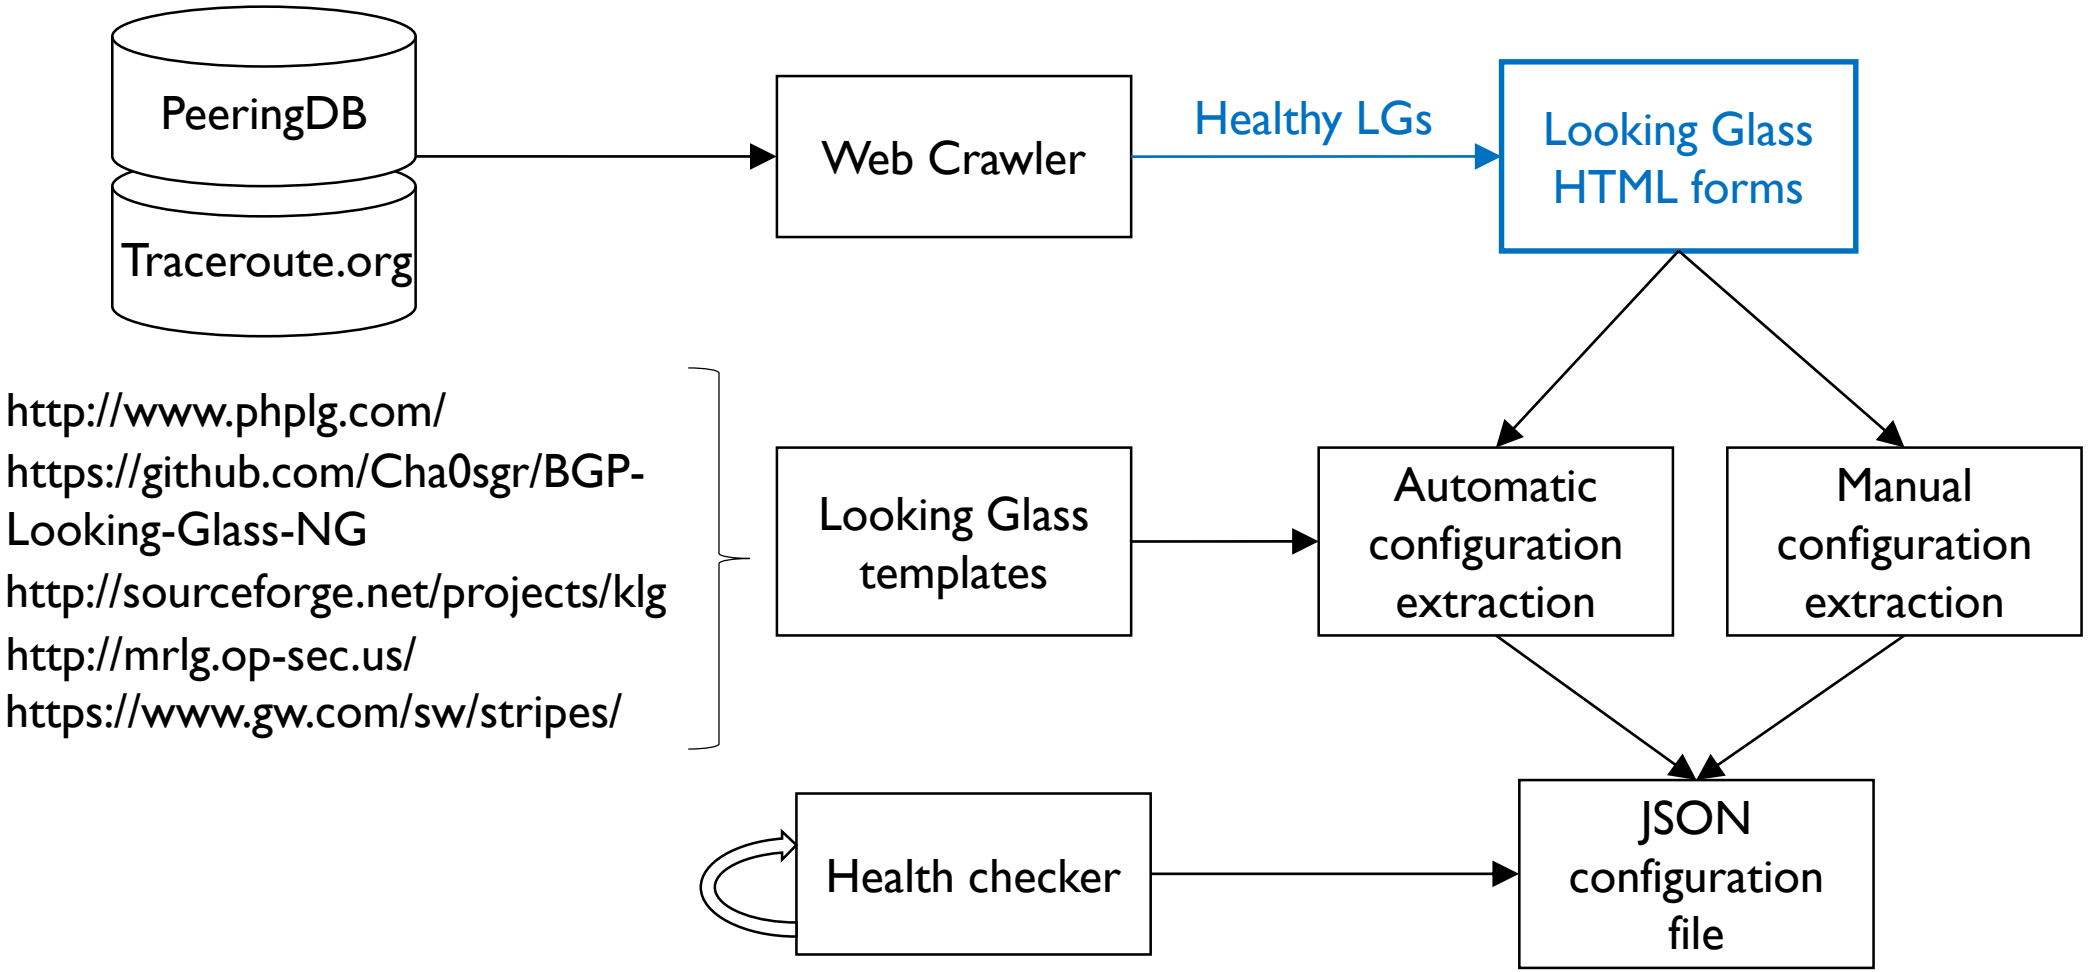

Looking Glass discovery workflow

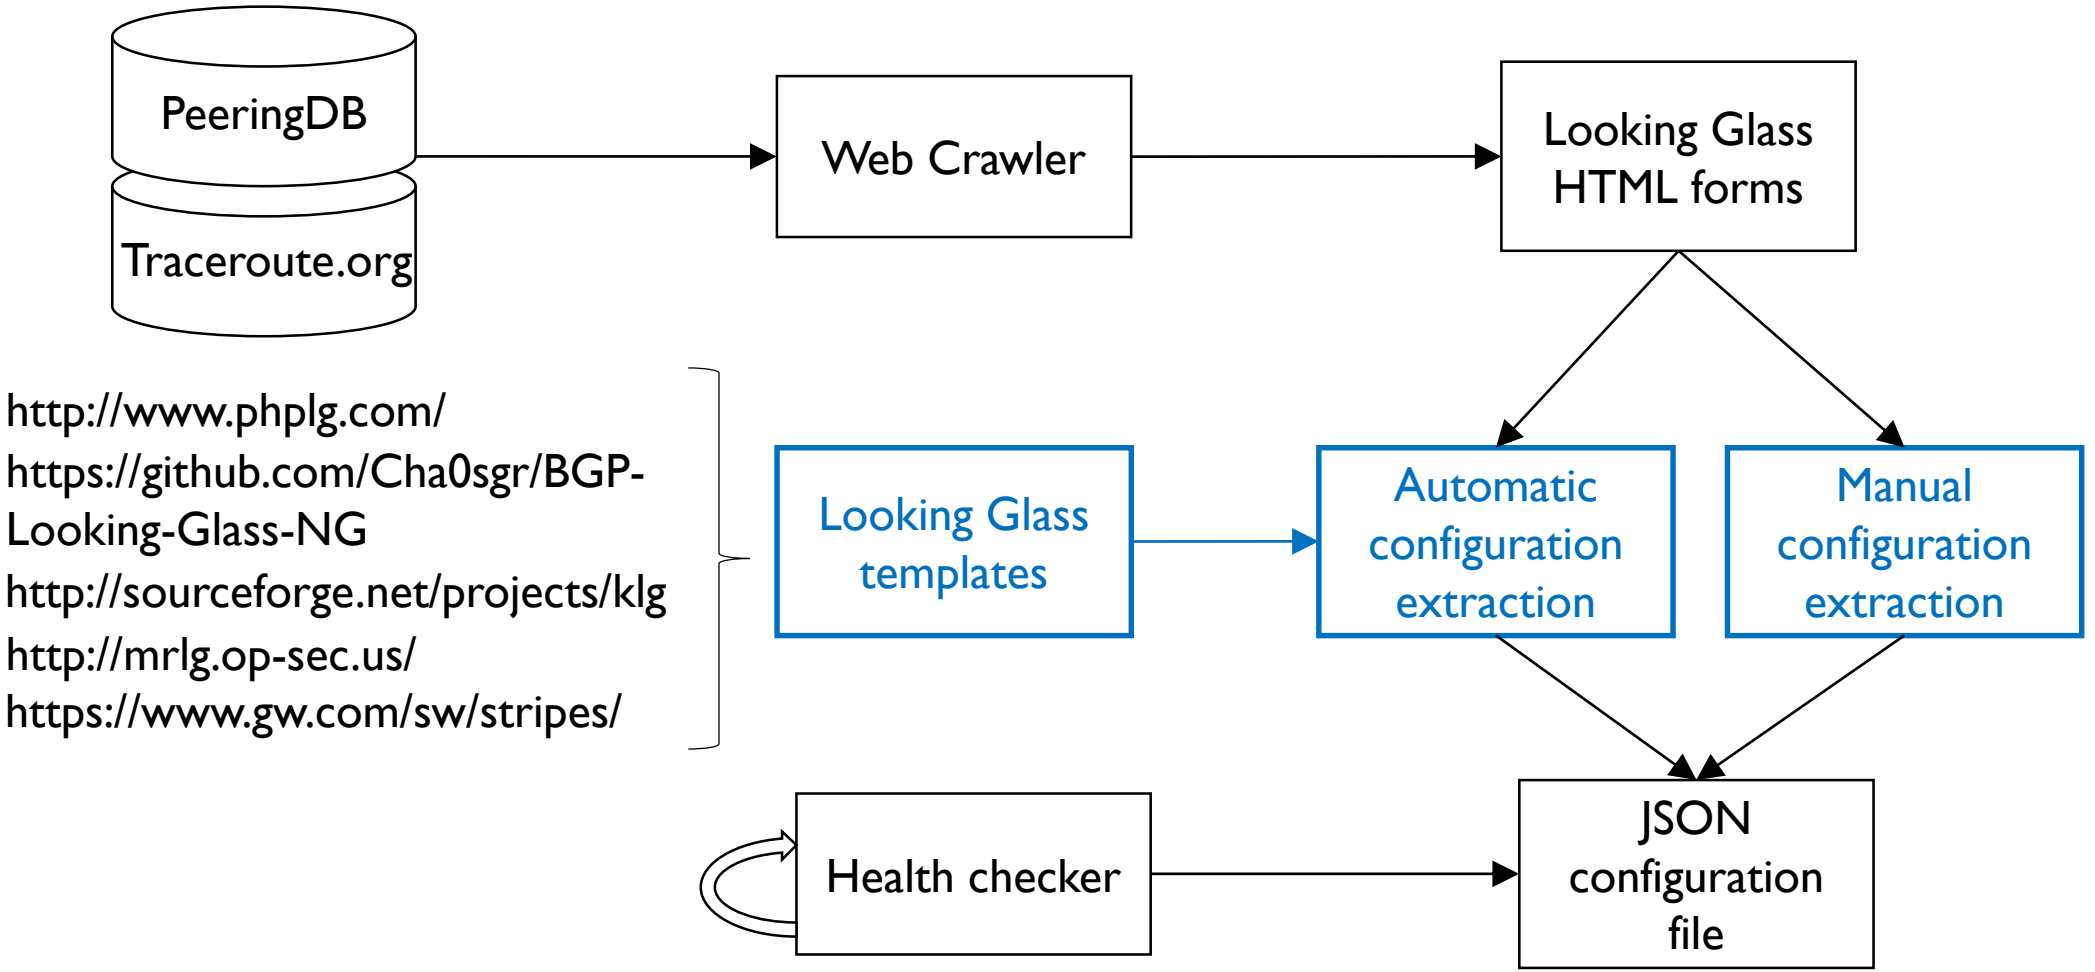

Looking Glass discovery workflow

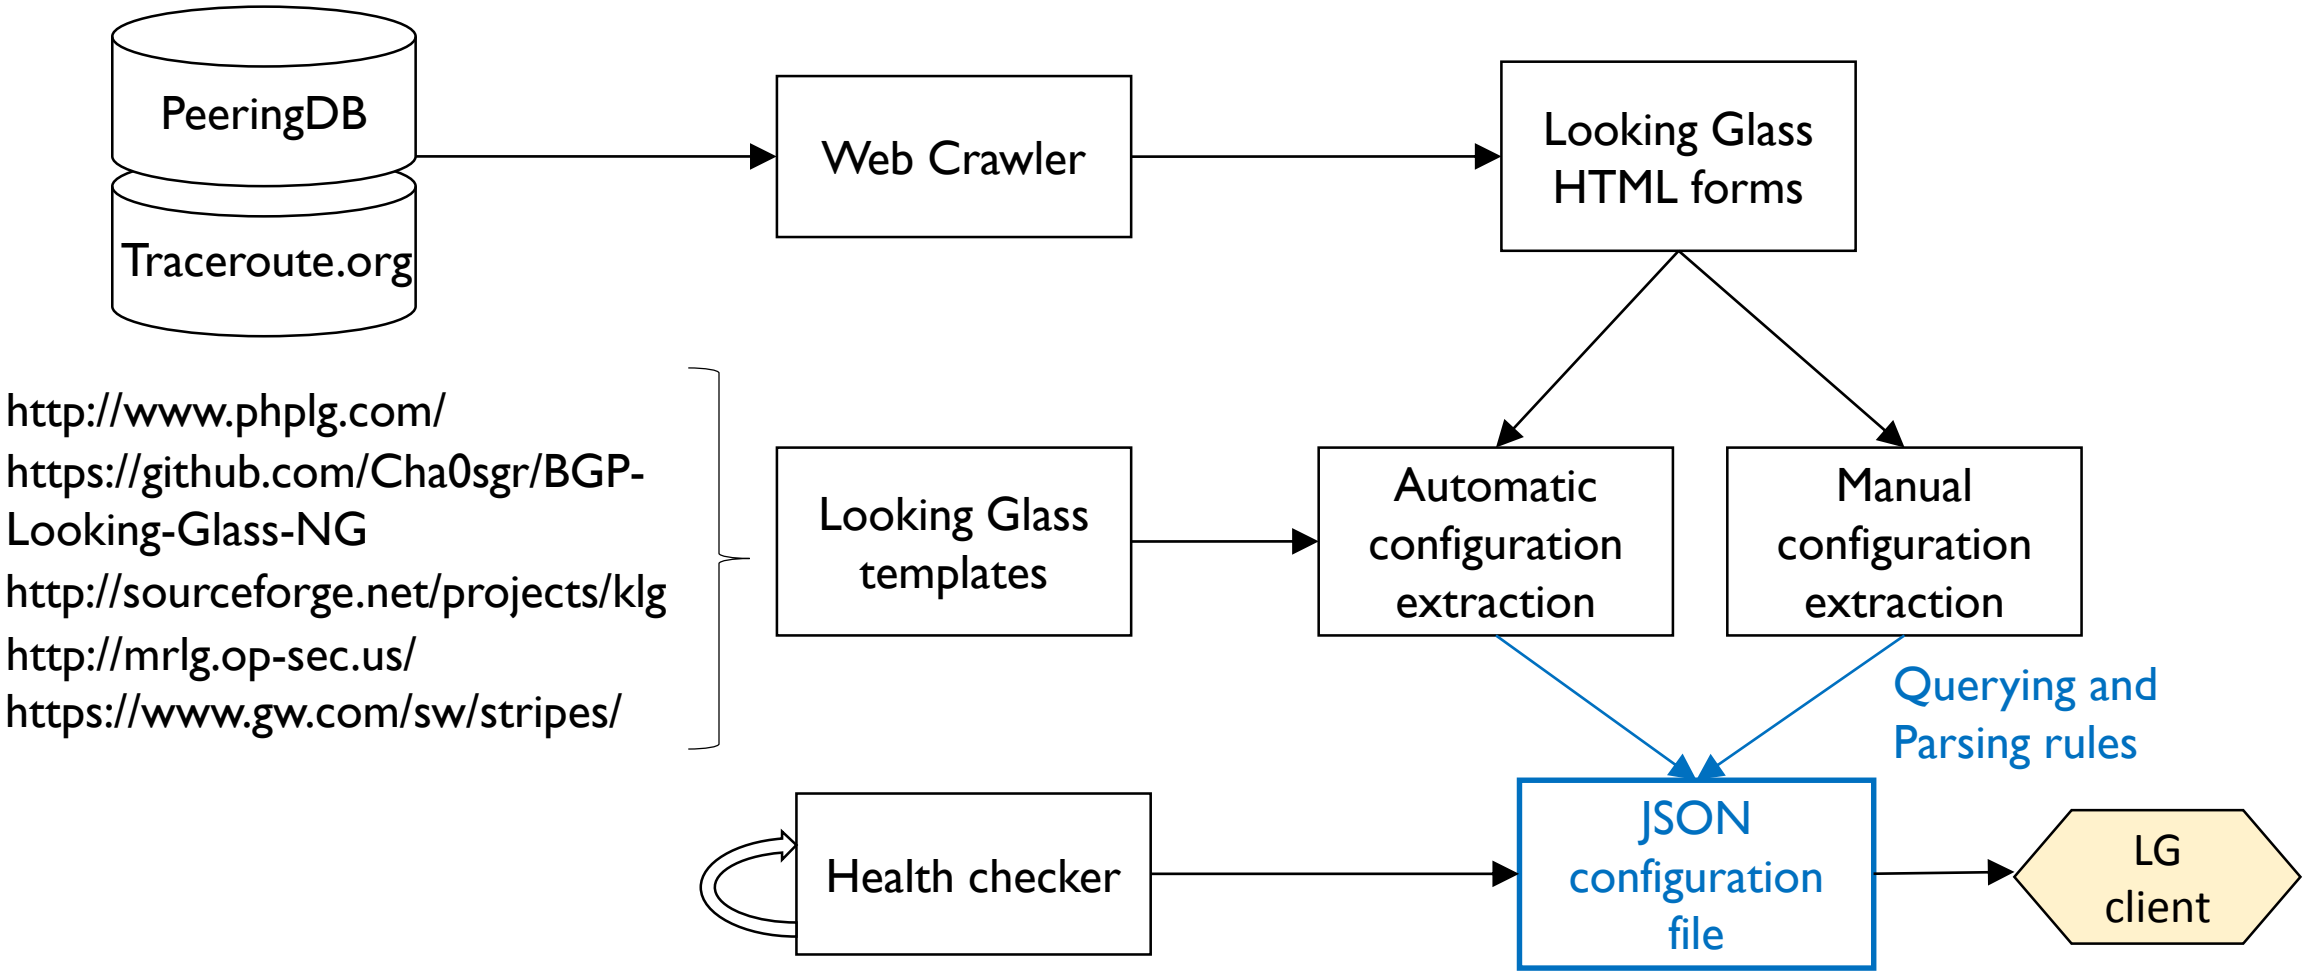

Looking Glass discovery workflow

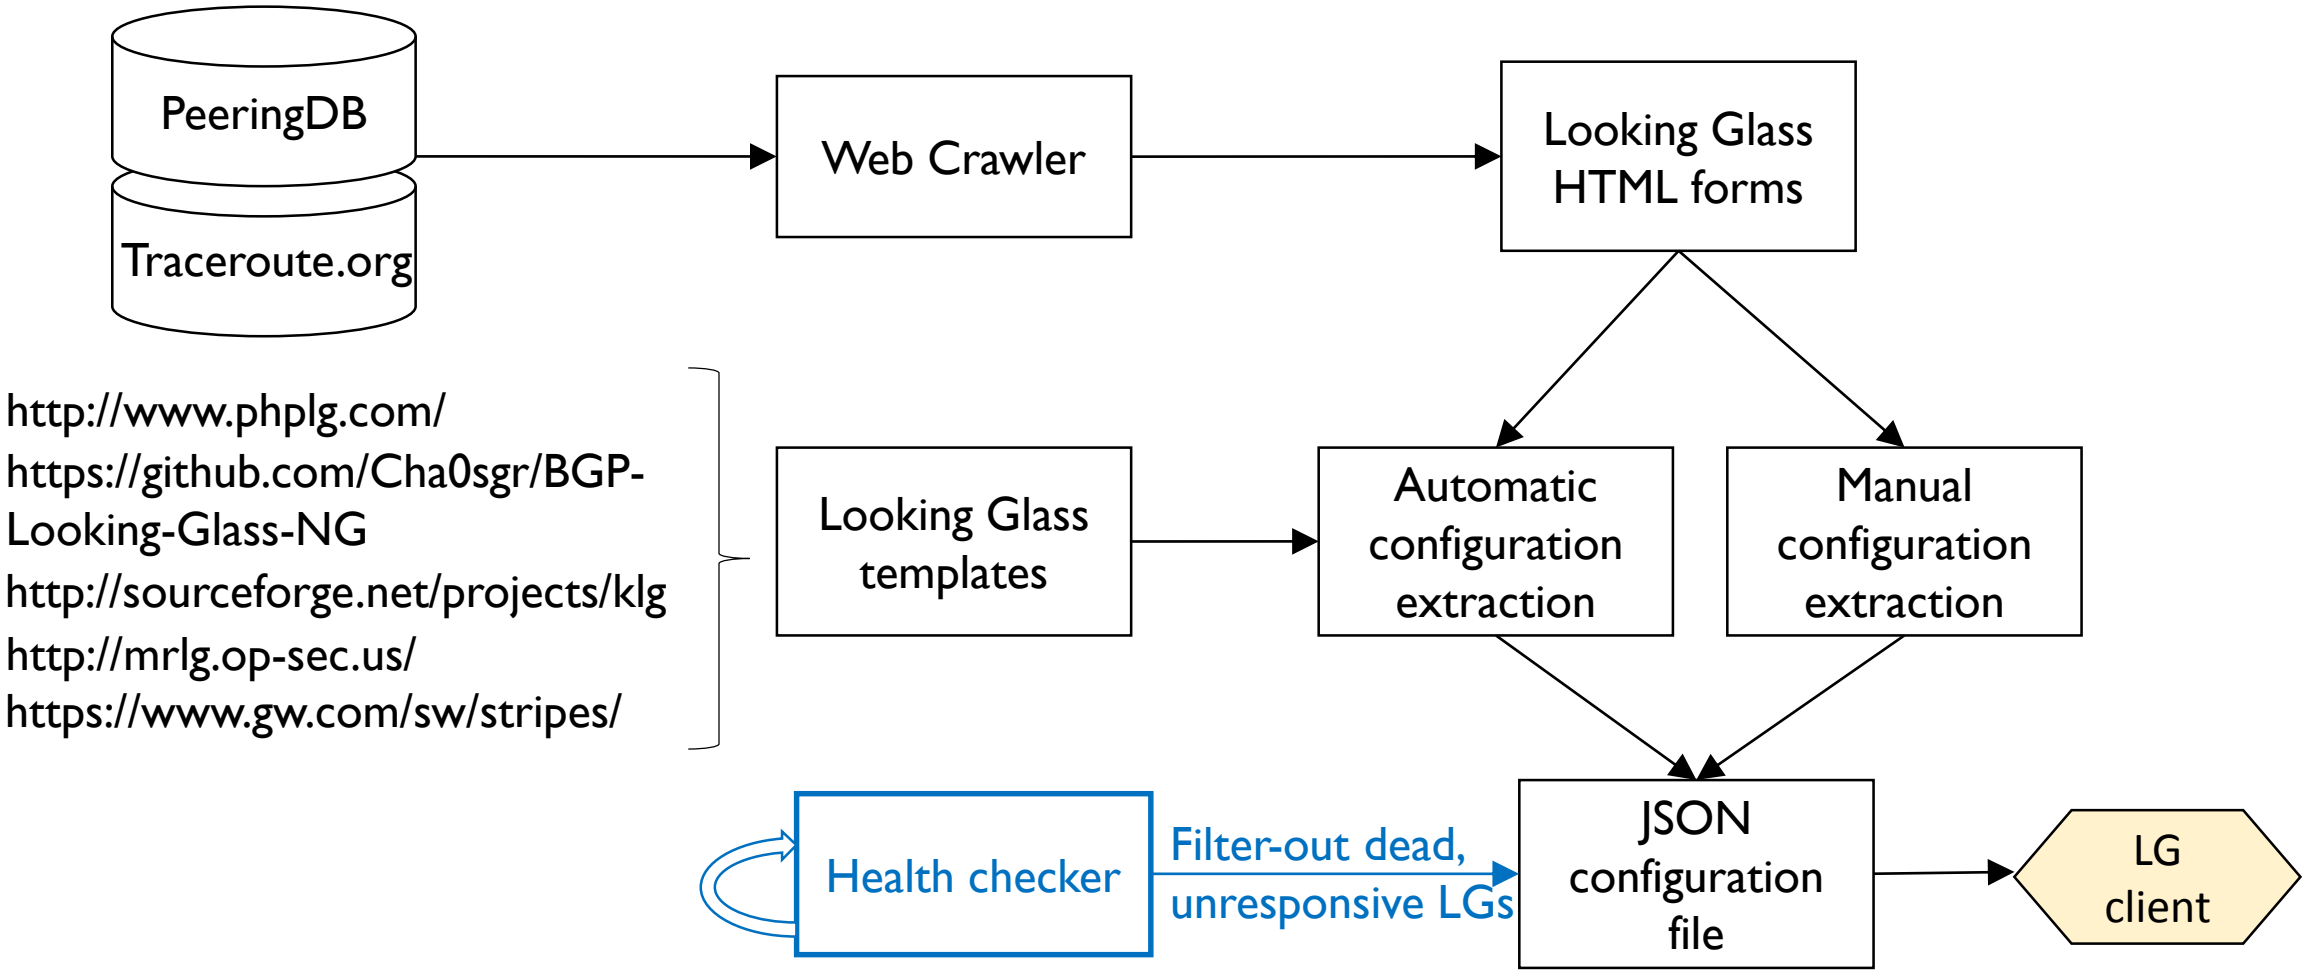

Support for multiple concurrent users requires multiple LG clients

Native LG querying

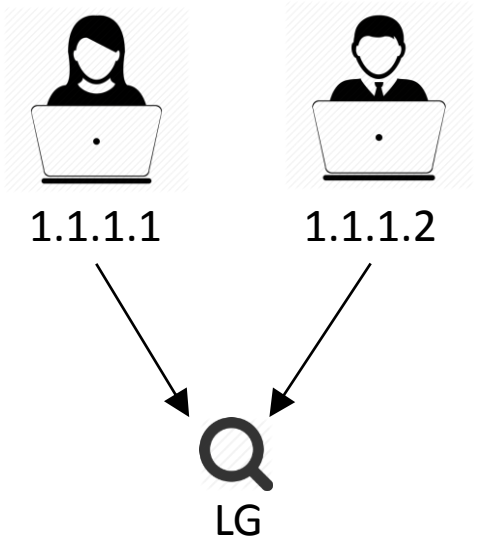

LGs use the users' IP address to impose per-user querying quotas

✗ Single-client Periscope

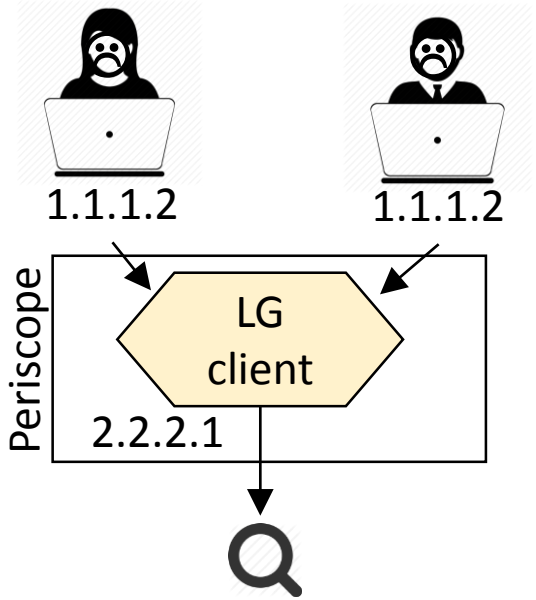

Putting multiple Periscope users behind the same IP causes all the users to share the quotas of a single user

✓ Multi-client Periscope

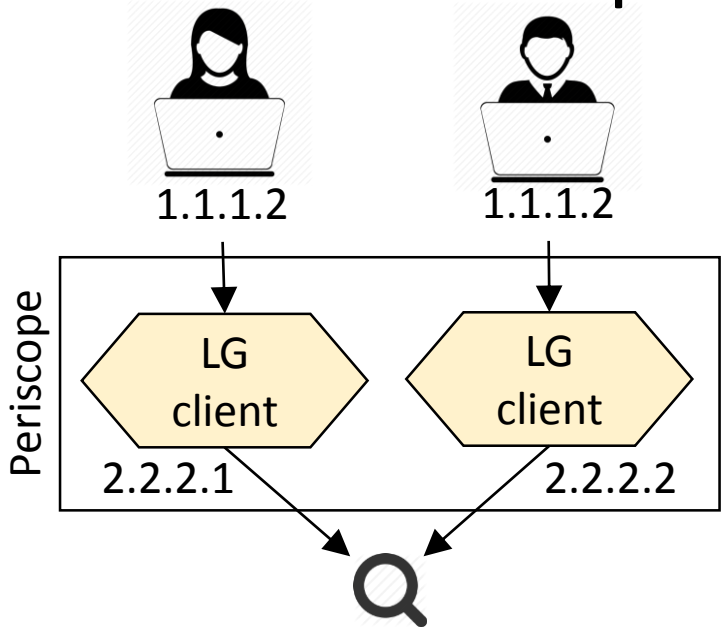

Using different client per user allows Periscope to provide the same querying quotas as native querying

Periscope's querying architecture

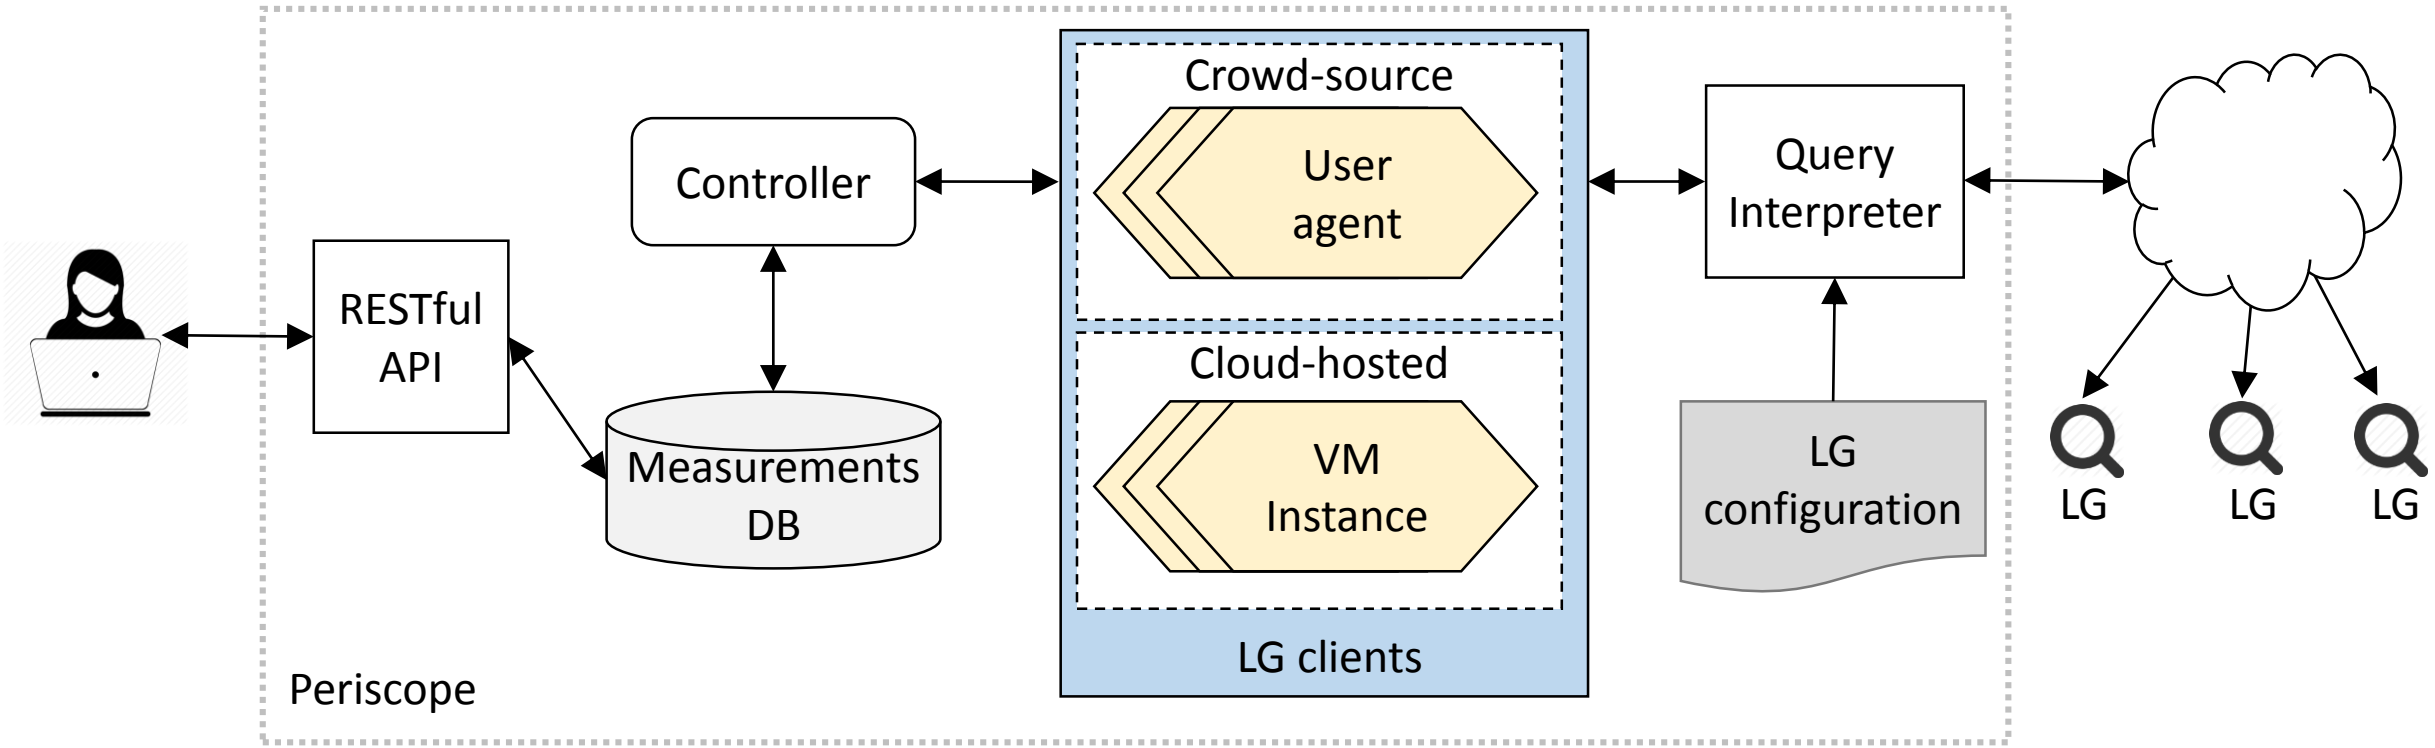

Abuse prevention

- The controller imposes two limits on the frequency of queries that can be submitted to a LG:
  - A *timeout* that expresses the minimum time interval between two consecutive LG queries by the same user
  - A number of *query slots* that indicate the maximum number of queries that Periscope will accept for an LG at any given moment
- If an LG has no available query slots it cannot be queried even if a user has not queried this LG for a period longer than the timeout

Geographical footprint of Periscope LGs

- As of December 2015 Periscope include LGs in **297** ASNs that provided **1,691** VPs in **76** countries
- Majority of VPs are located in Europe and the US

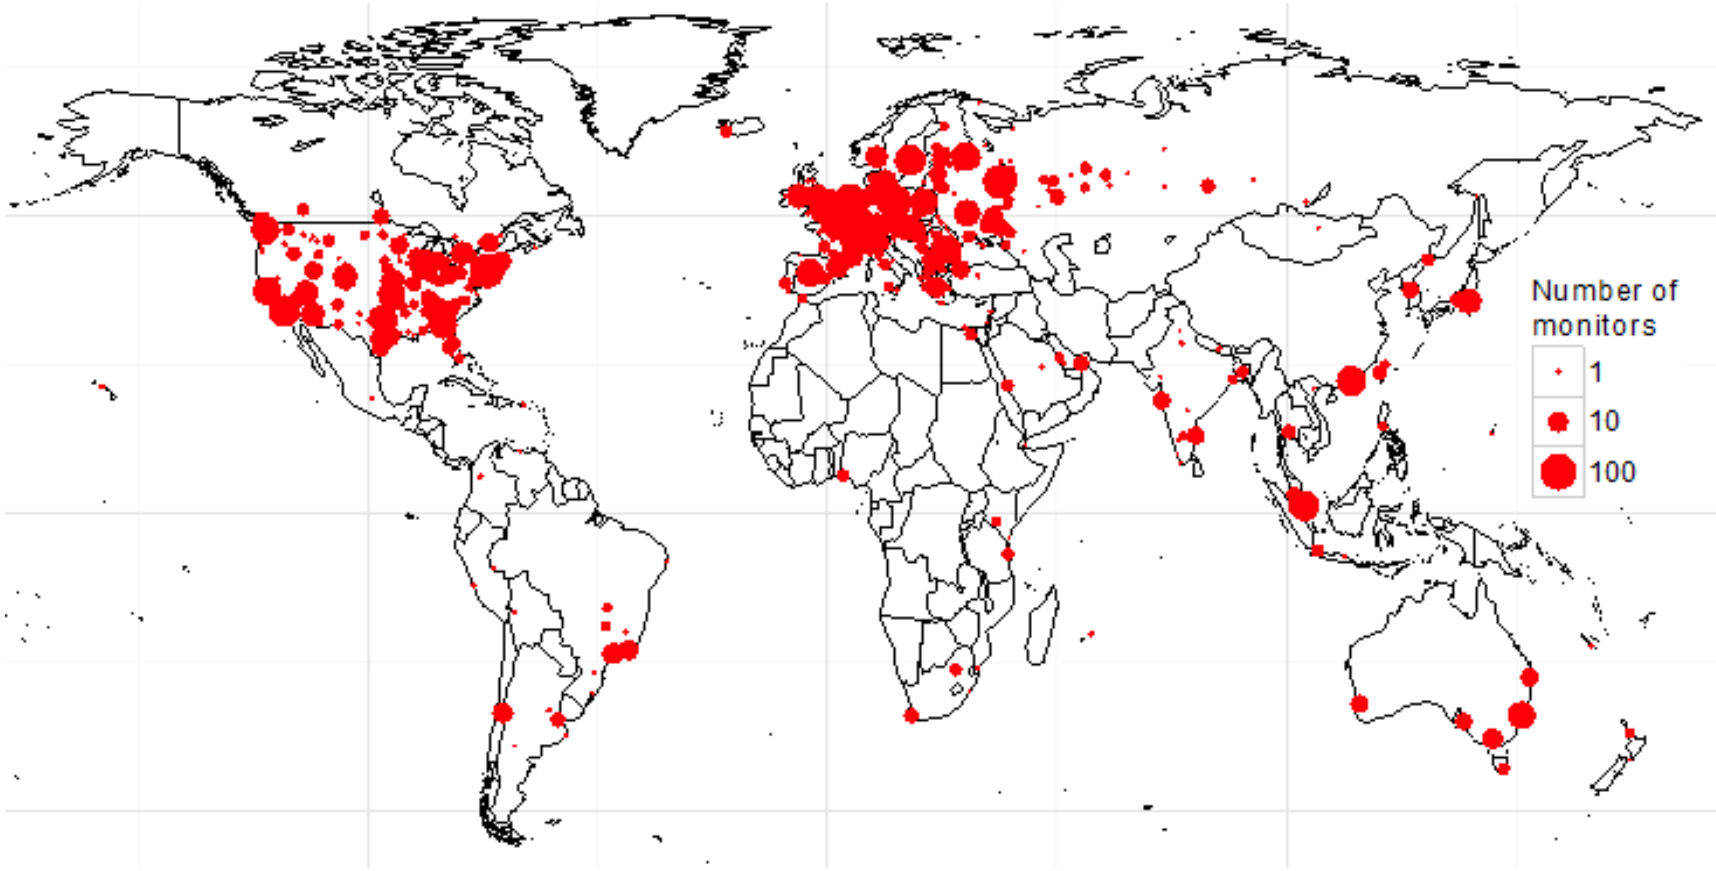

Commands supported by Periscope LGs

- Over 75% of the LG nodes provide both traceroute and BGP commands
- Over 60% of the LGs support IPv6 queries

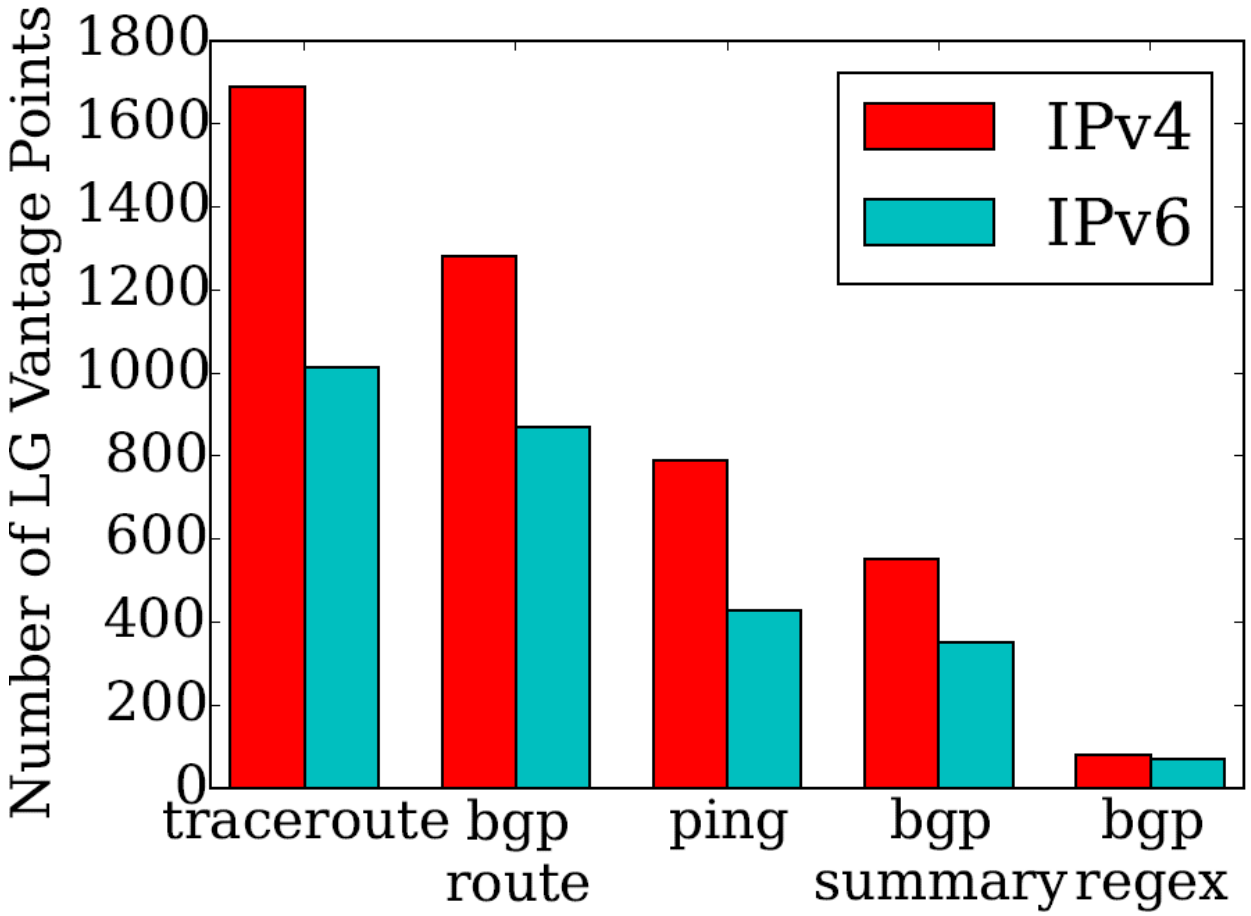

The topology observed by LGs is largely complementary to other platforms

Datasets	ASes		AS links		IXPs	
	Observed	Unique	Observed	Unique	Observed	Unique
LG	3109	809 (17%)	29525	13969 (19%)	167	16 (8%)
Atlas *	3369	1464 (31%)	55936	40620 (55%)	171	21(10%)
Ark	1608	59 (1.2%)	10237	1625 (2.2%)	136	8 (4%)
All	4657	-	73348	-	202	-

- Queried 2,000 randomly selected IPs from each LG and from each VP available in RIPE Atlas and CAIDA's Ark

* RIPE Atlas measurements were executed using an account with elevated probing quota that allowed the collection of 6 million traceroutes from the all the active probes in a period of 2 months

The unique ASes in each dataset differ significantly in terms of customer cone sizes.

- The customer cone is the number of ASes in the downstream path of a given AS
- LGs tend to capture more peripheral and stub ASes, while Ark and Atlas capture ASes with larger customer cone.

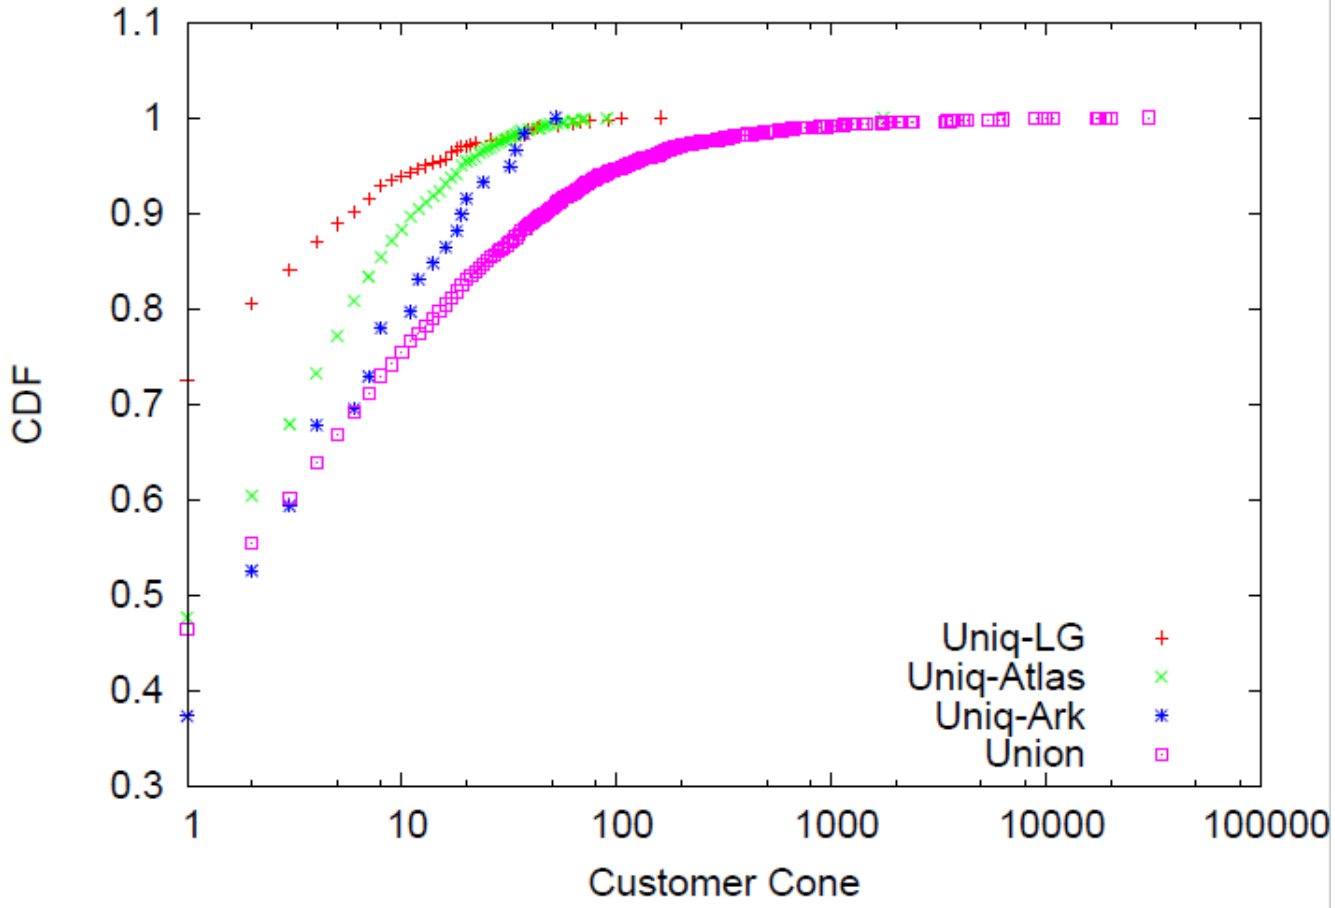

Case studies

Validation of IP-to-AS mapping

- Validation of IP-to-AS mapping requires comparison of BGP and traceroute paths obtained from VPs as close as possible inside the same AS to avoid path mismatches due to intra-domain routing dynamics

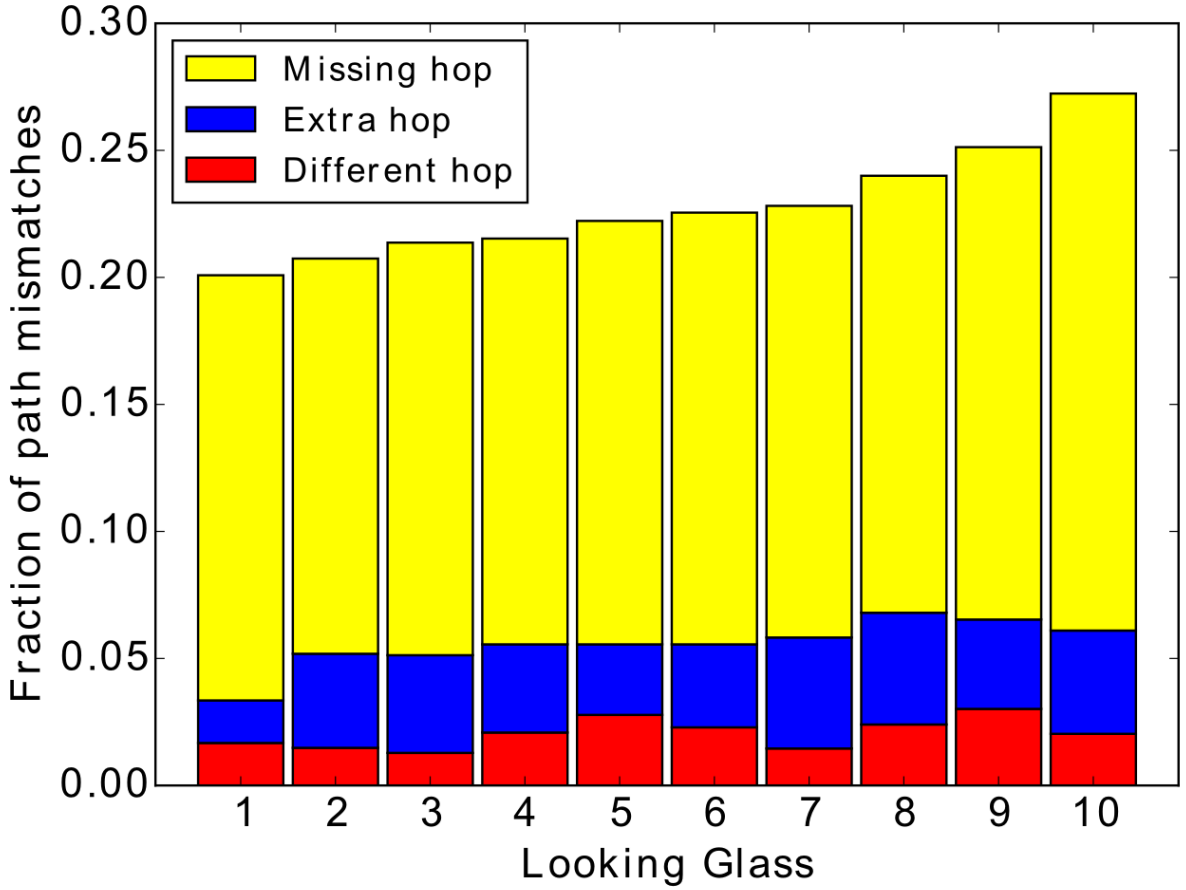

Geolocation of IP interfaces of border routers

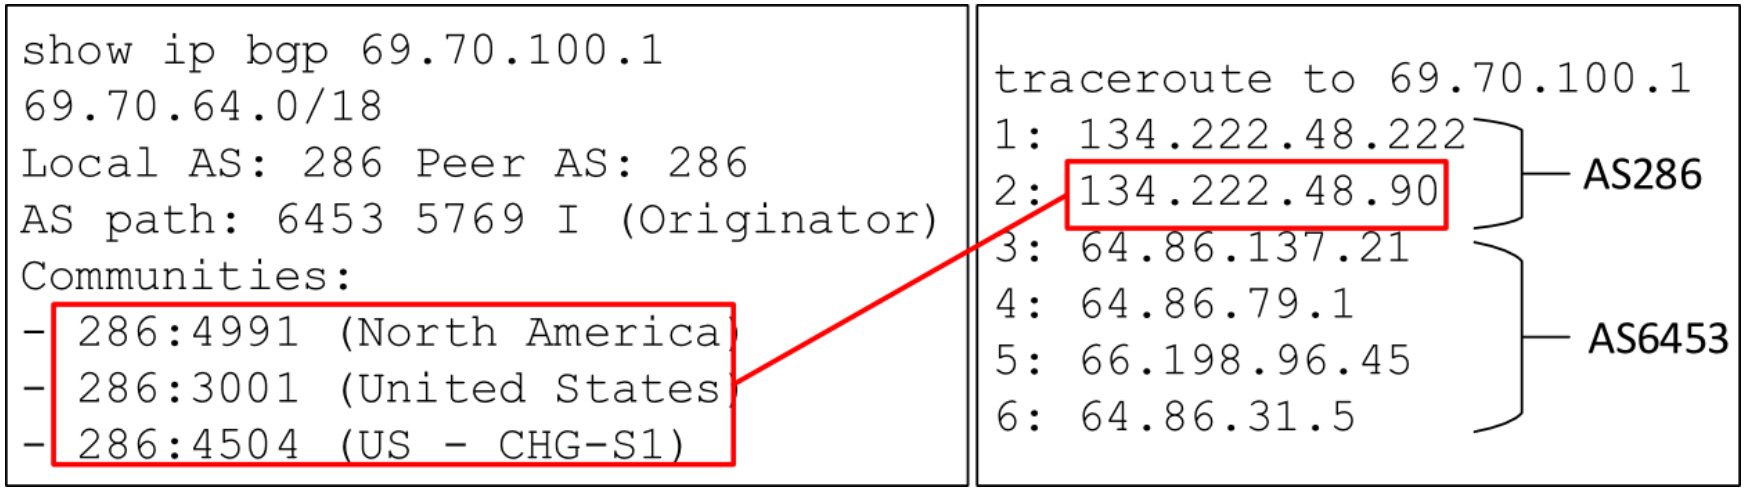

- Operators often use BGP communities to tag the entry point of a route
- Combine BGP communities and traceroute paths from the same VP to associate locations encoded in communities with router interfaces

Conclusion and ongoing work

- Access to Periscope can be provided after email request: vgiotsas@caida.org
- API documentation: http://www.caida.org/tools/utilities/looking-glass-api/
- Ongoing work:
  - Development of a user interface
  - Integration with RIPE Atlas, Ark and BGPStream to improve orchestration of measurement campaigns

Backup slides

HTML Template for the input parameters of the Version6 LG

Input name	Input type	Expected values	Meaning
query	radio	[bgp, trace, ping]	[show ip bgp, traceroute, ping]
addr	text	*	Query target
router	select	*	Router identifier
protocol	select	[IPv4, IPv6]	IP version