PERISCOPE: Standardizing and Orchestrating Looking Glass Querying

Vasileios Giotsas
vgiotsas@caida.org

Purpose of this talk

- Inform the operational community about Periscope.
- Solicit feedback:
  - Details that we may have missed
  - Ways to make Periscope more useful
  - Technical insights, usage statistics, historical data …
- Encourage engagement and contributions

High-level goals and principles of Periscope

Periscope unifies the discovery and querying of Looking Glasses under a uniform API

- Respect resource limitations and preserve conservative query rates
- Provide transparency and accountability in Looking Glass querying.
- Be responsive and compliant to operators' requests.

Advantages of LG measurements

- LGs are among the few public measurement tools that provide direct interfaces to routers:
  - Access to non-transitive BGP attributes (e.g. LocPref).
  - Co-located BGP and traceroute/ping monitors.
  - Vantage Points at colocation facilities, IXPs, datacenters.
  - Vantage Points in ASes not covered by other platforms

- Motamedi, R., Rejaie, R., & Willinger, W. (2015). A Survey of Techniques for Internet Topology Discovery. *Communications Surveys & Tutorials, IEEE*, *17*(2), 1044-1065.
- A. Khan, T. Kwon, H.-c. Kim, & Y. Choi, “AS-level Topology Collection Through Looking Glass Servers,” in IMC ’13

Problems with LGs

- Lack of standardization and consistency:
  - Disparate interfaces and output formats
- Hard to discover and track:
  - No centralized index of LGs, their locations and their capabilities
- High attrition rates:
  - Hard to maintain an up-to-date list of LGs

M. Stubbig, “Looking Glass API.” https://tools.ietf.org/html/draft-mst-lgapi-04, May 2016.

Problems with LGs

- Lack of standardization and consistency:
  - Disparate interfaces and output formats

Periscope implements a common querying scheme, indexing and data persistence features

- High attrition rates:
  - Hard to maintain an up-to-date list of LGs

M. Stubbig, “Looking Glass API.” https://tools.ietf.org/html/draft-mst-lgapi-04, May 2016.

Periscope Workflow

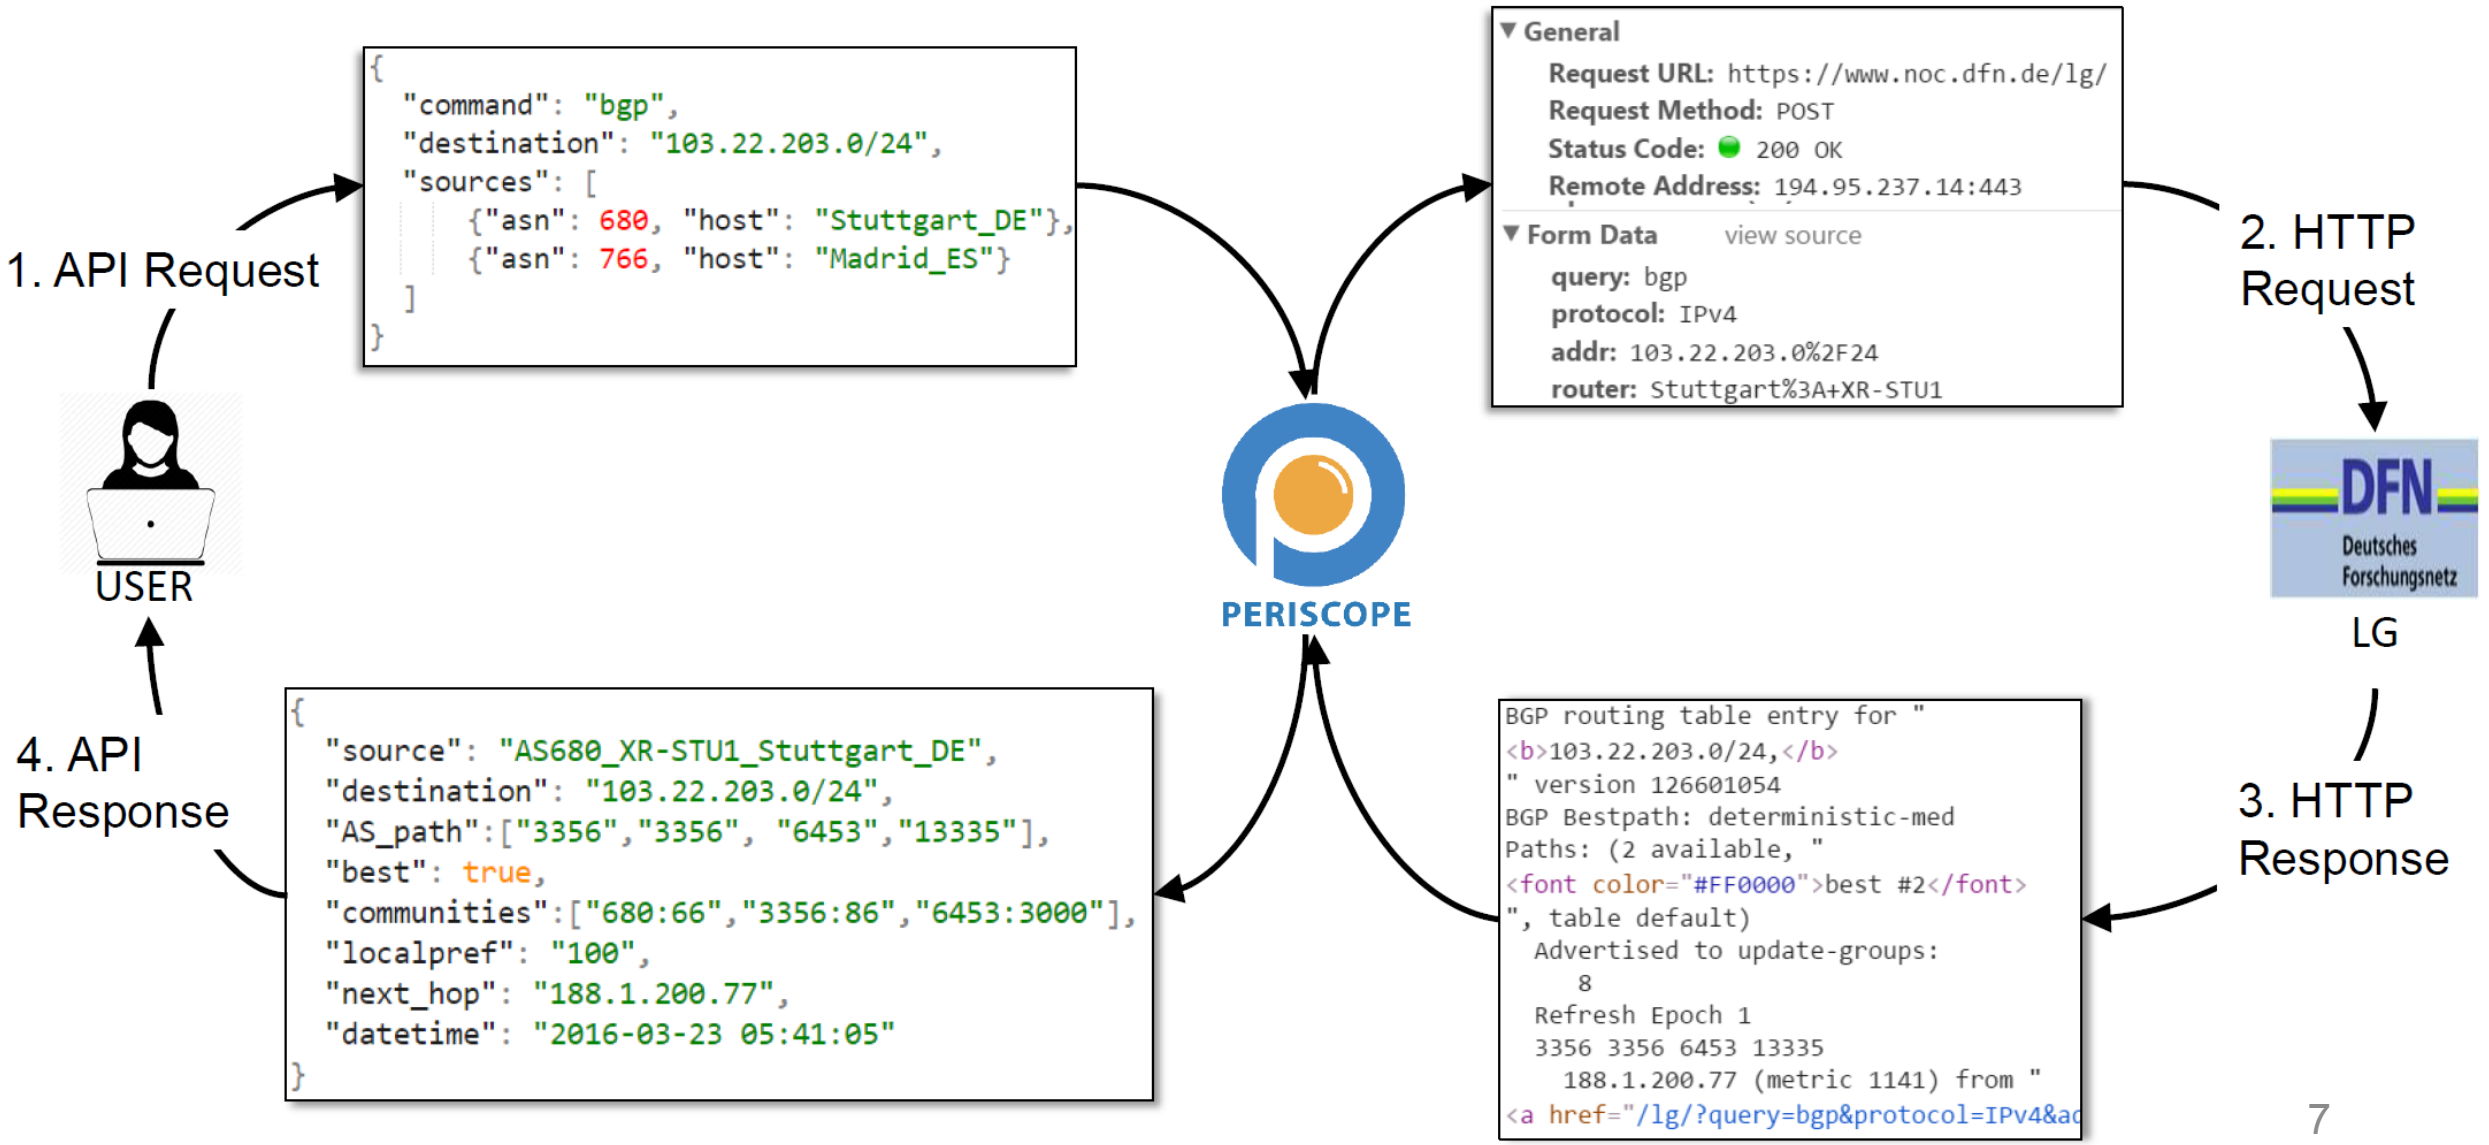

Periscope Architecture

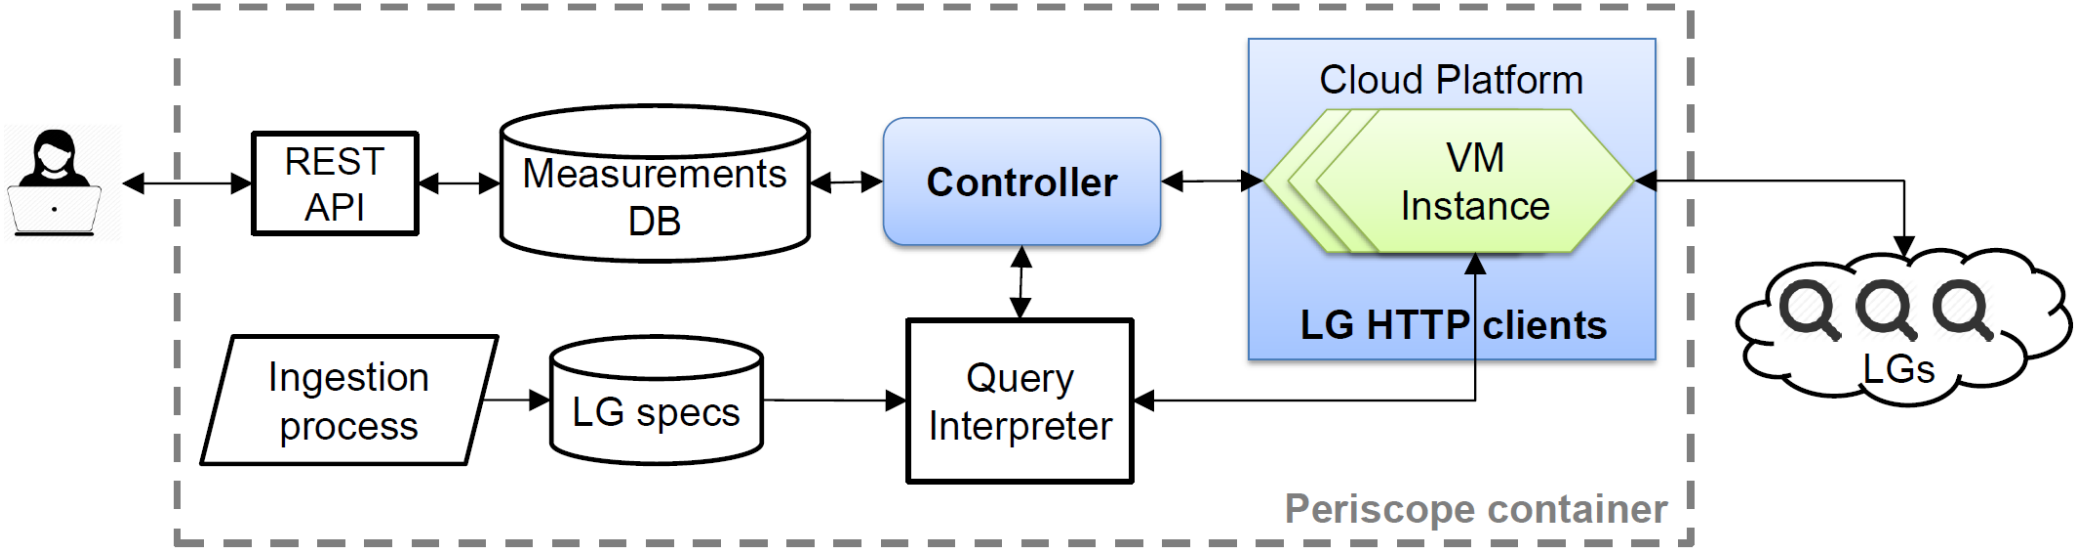

- For each Periscope User the controller allocates a different cloud-hosted **VM instance** to execute the user queries.
- Each VM instance takes an IP address from the cloud operator's address space.
- The controller implements throttling of query rates.

Periscope enforces per-user and per-LG query rate limits

- Two limits control the rate of issued LG queries:
  - **User-specific**: Each user can issue only 1 query per 5 minutes to the same LG.
  - **LG-specific**: Each LG will execute up to 3 queries per minute from all the users.
- A query is allocated if neither limit is exceeded.
- Exponential back off when LGs respond with errors

Transparency and accountability

- Periscope sets three custom HTTP headers in every request:
  - **"X-Request-Origin: periscope"**
  - **"X-Request-For:<*user-ip*>"**
  - **"X-Request-Client:<*gcloud* OR *aws* OR *ark*>"**

- Periscope IPs configured with reverse DNS records.
- Periscope assigns an LG client with a *static* IP address to each Periscope user to allow persistence identification.

Coverage of Periscope LGs

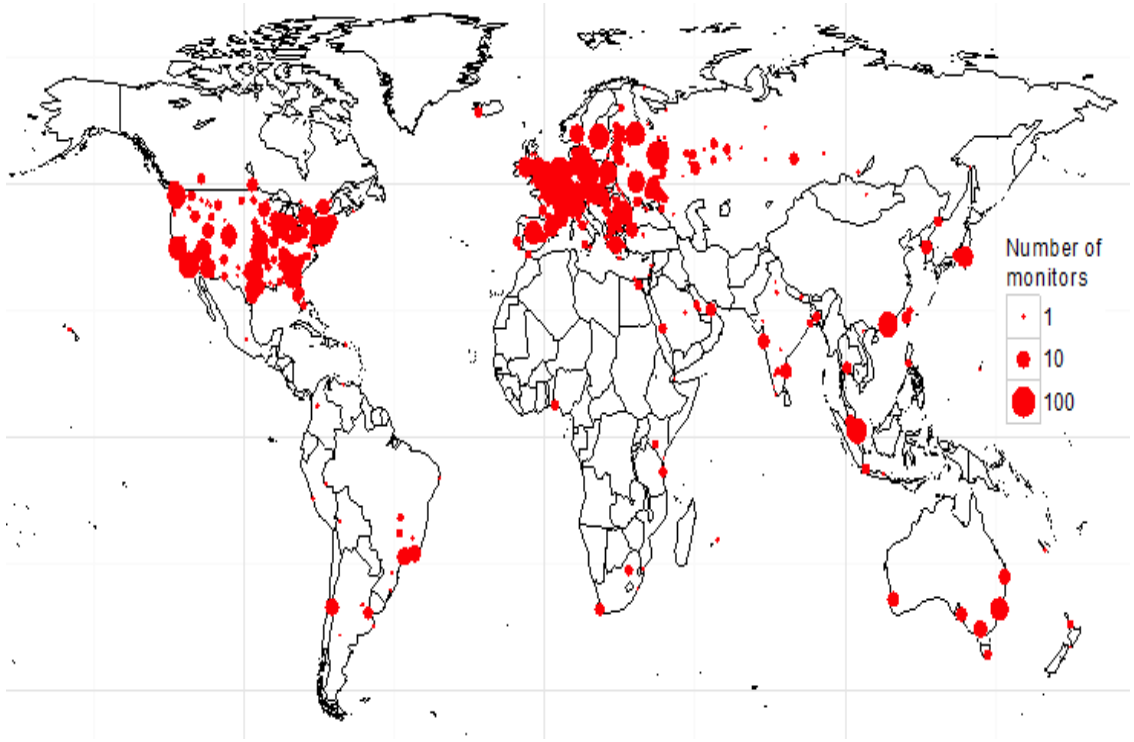

- 572 ASNs with 2,951 VPs.
- 77 countries, 492 cities.

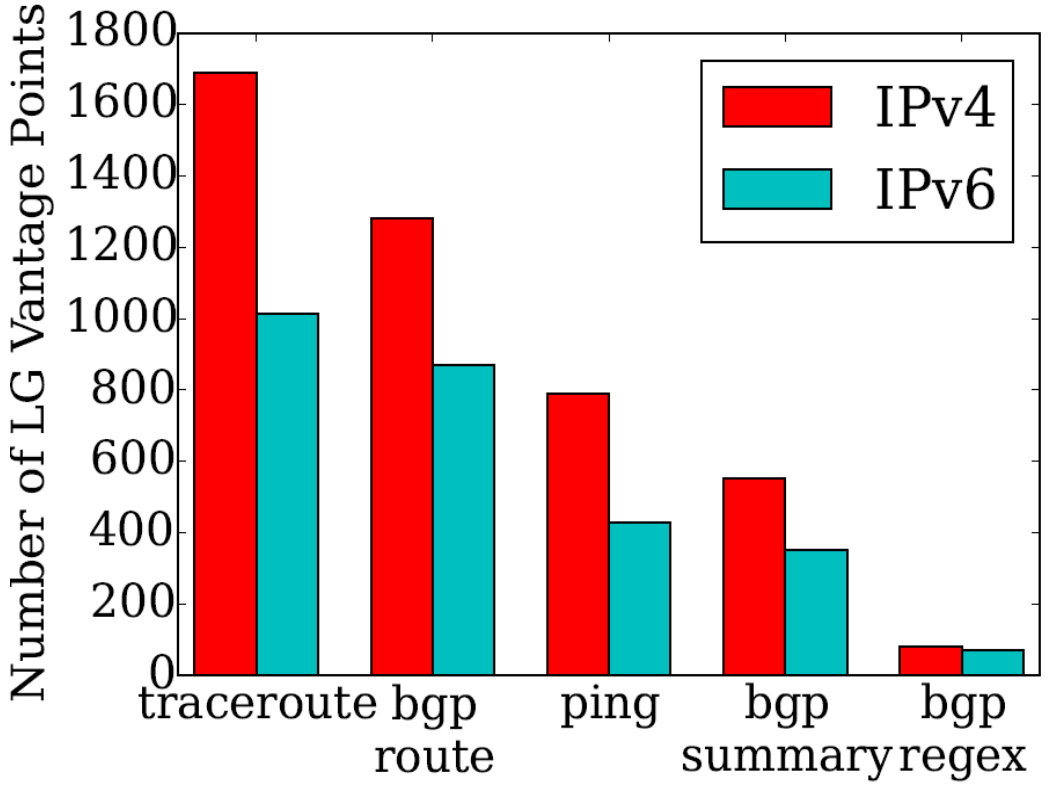

- > 75% of VPs provide both traceroute and BGP.
- > 60% of LGs support IPv6.

The topology observed by LGs is largely complementary to other platforms

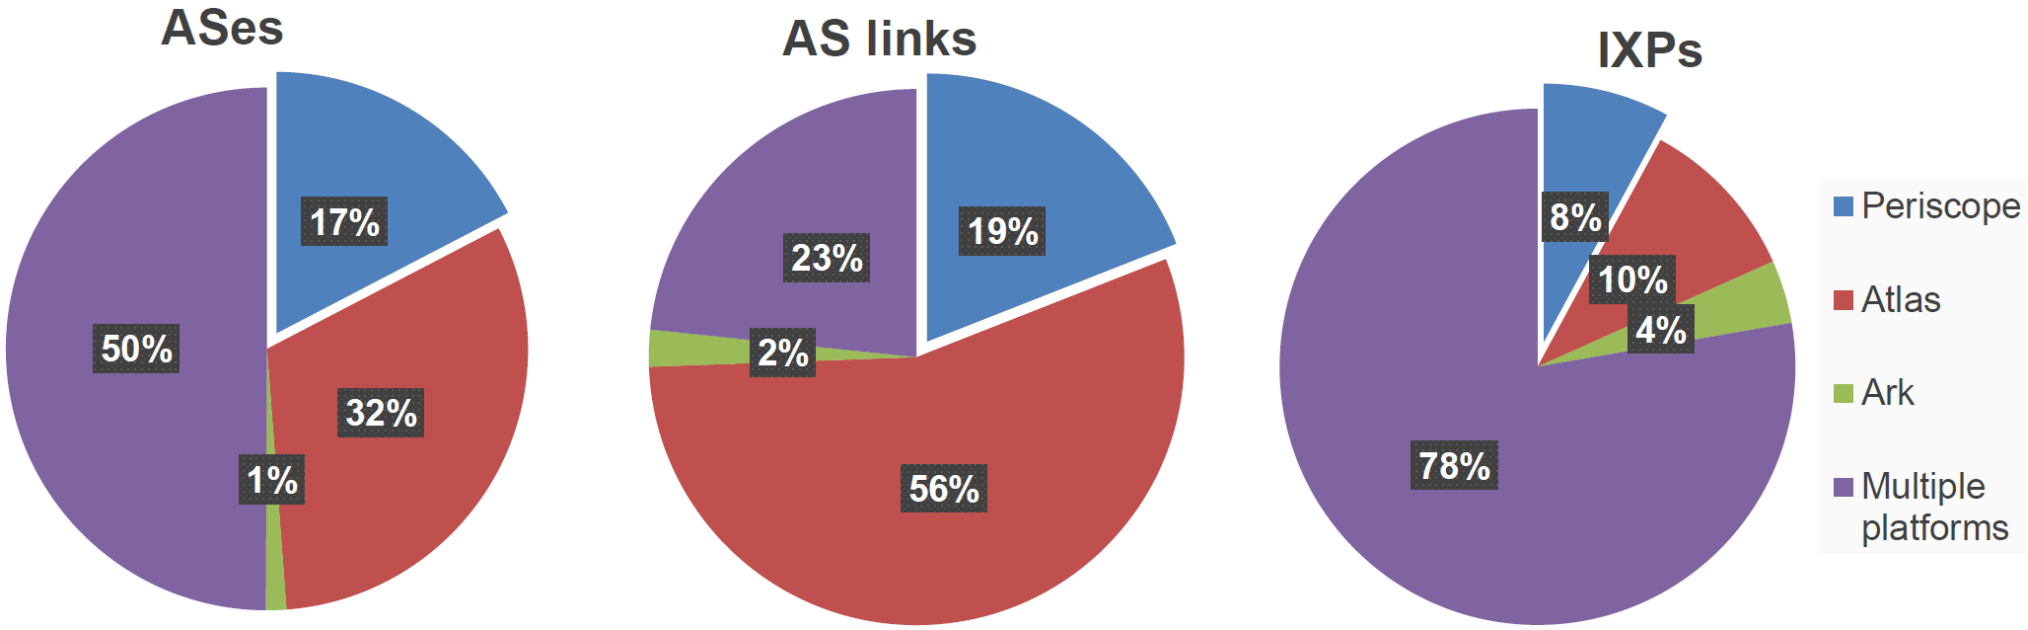

- Queried 2,000 randomly selected IPs from each LG and from each VP available in RIPE Atlas and CAIDA's Ark

Benefits

- Easier to discover and query new VPs for reverse paths
- Easier policing of Looking Glass usage through an access-control layer
- Improved utilization and load distribution
- Avoid redundant measurements by archiving and making public historical measurement data.

Request for contributions

- Please contribute feedback regarding:
  1. Per user query limits
  2. Global query limits
  3. Opt-in or opt-out requests
- Utilization statistics and archived queries.
- Funding, infrastructure support (VM instances, cloud computing credit).

Conclusion

- Periscope goals:
  - Unify LGs under a uniform API.
  - Enforce per-user and global query limits.
  - Provide Transparency and accountability
- Access request: periscope-info@caida.org
- Documentation:

  http://www.caida.org/tools/utilities/looking-glass-api/

Questions?

PERISCOPE caida

BACKUP SLIDES

Support for multiple concurrent users requires multiple LG clients

Native LG querying

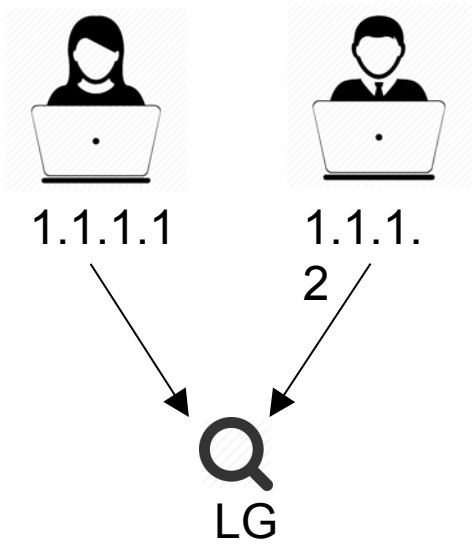

LGs use the users' IP address to impose per-user querying quotas

✘ Single-client Periscope

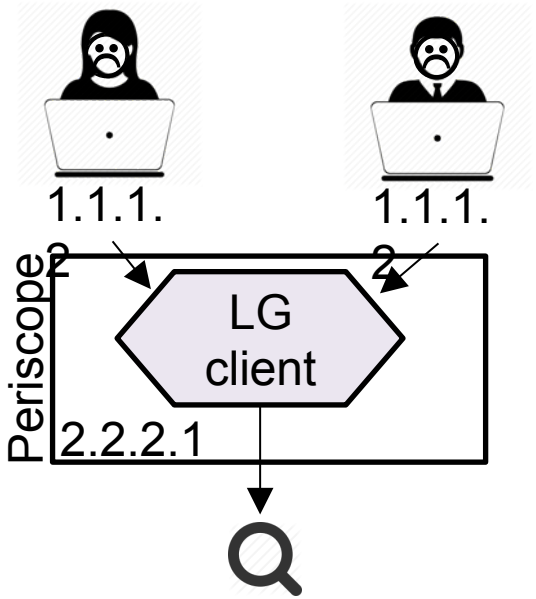

Putting multiple Periscope users behind the same IP causes all the users to share the quotas of a single user

✔Multi-client Periscope

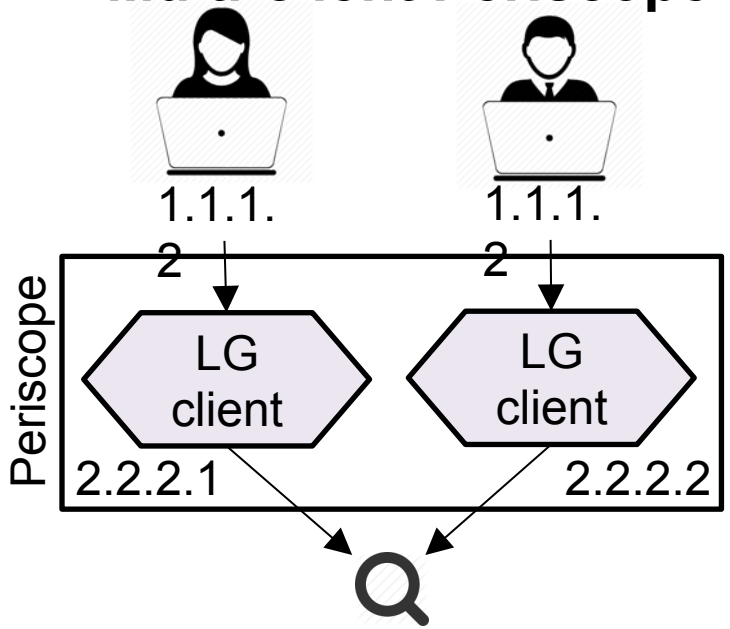

Using different client per user allows Periscope to provide the same querying quotas as native querying

LG Ingestion Workflow

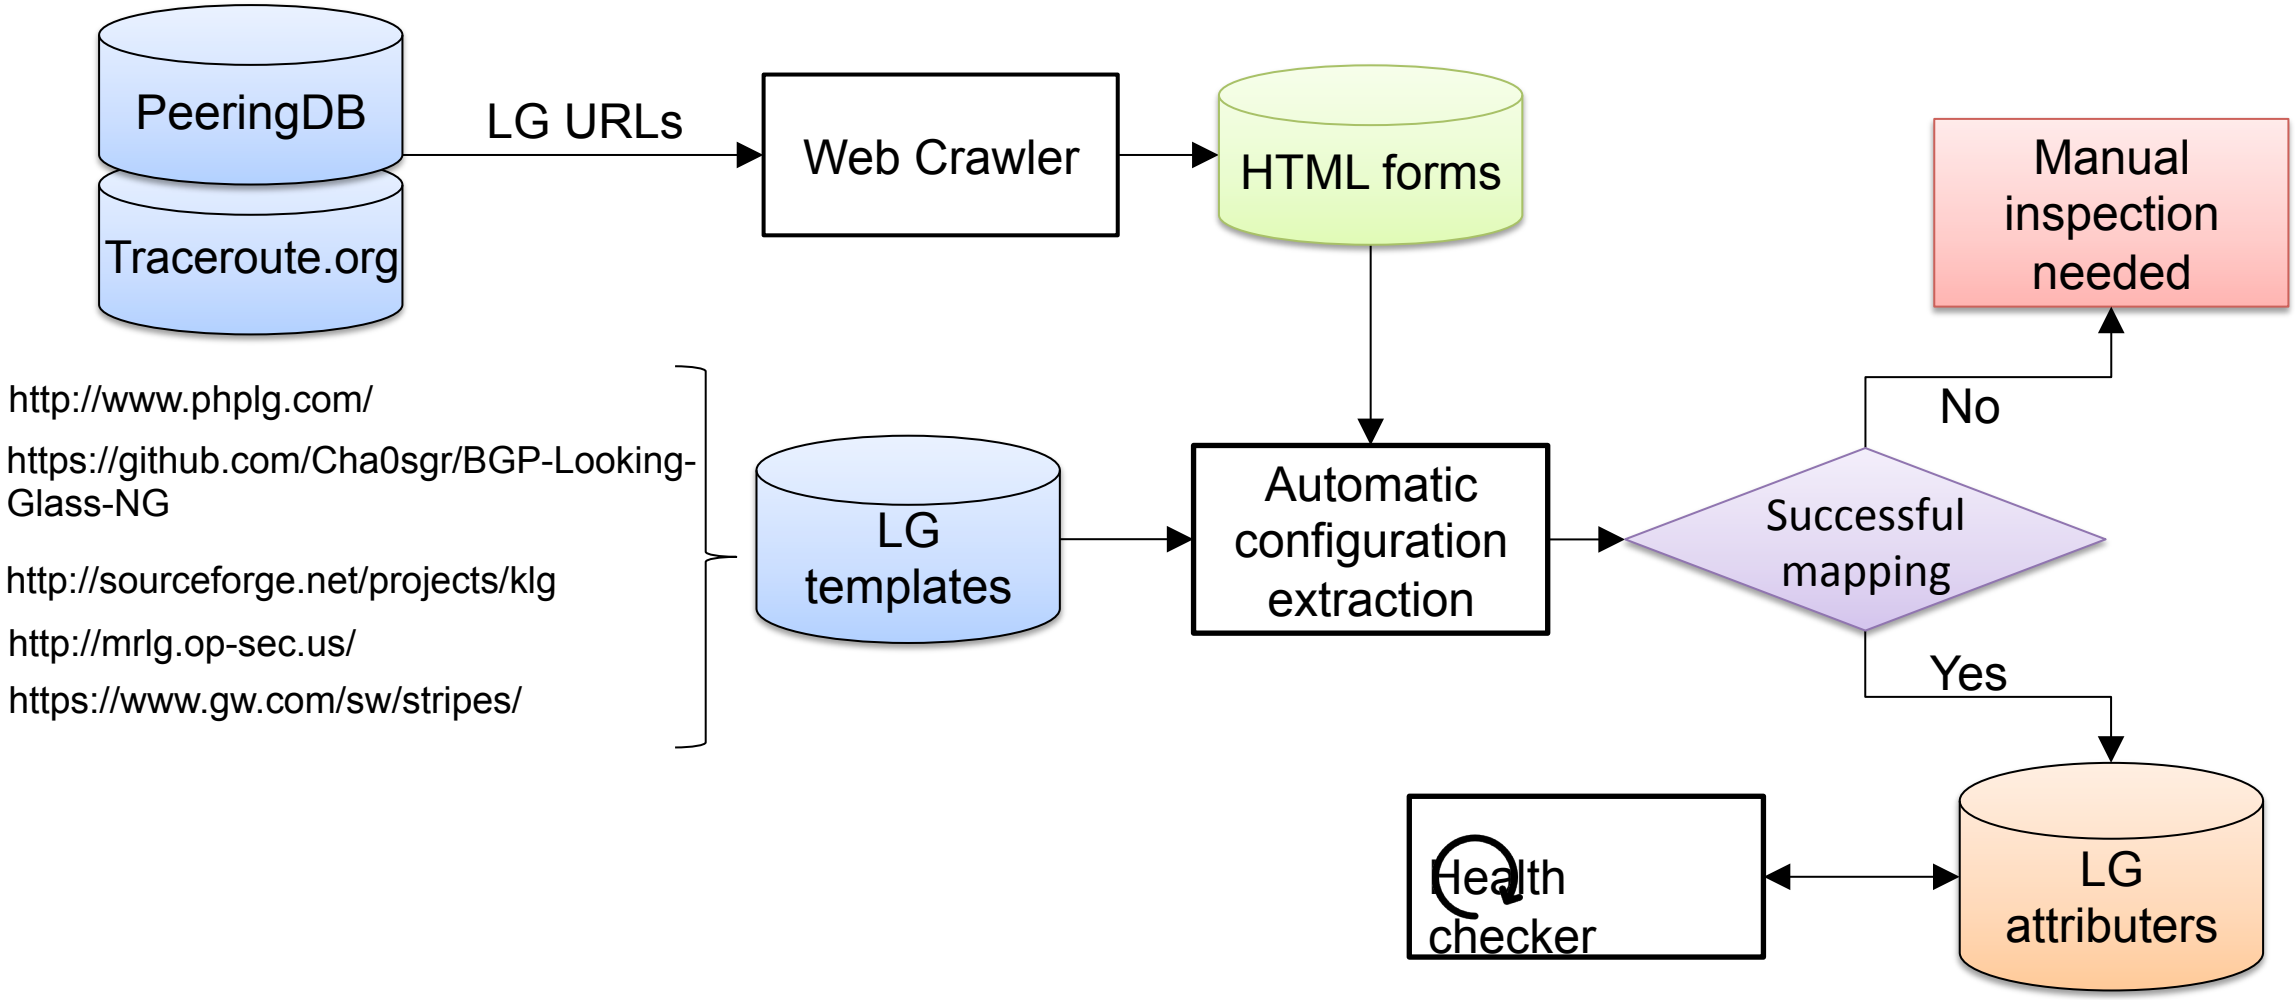

LG Ingestion Workflow

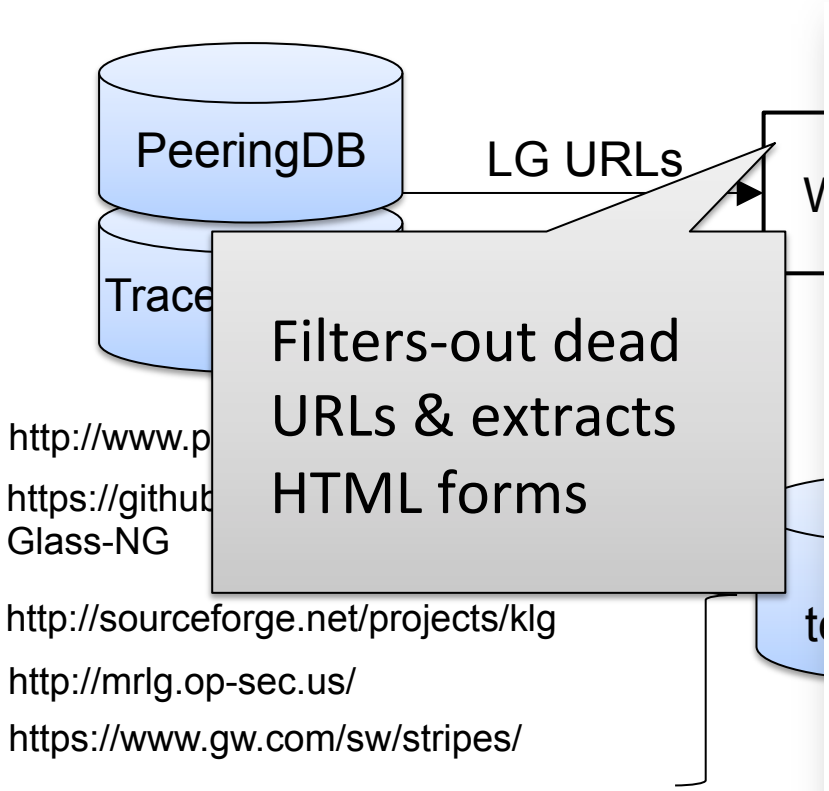

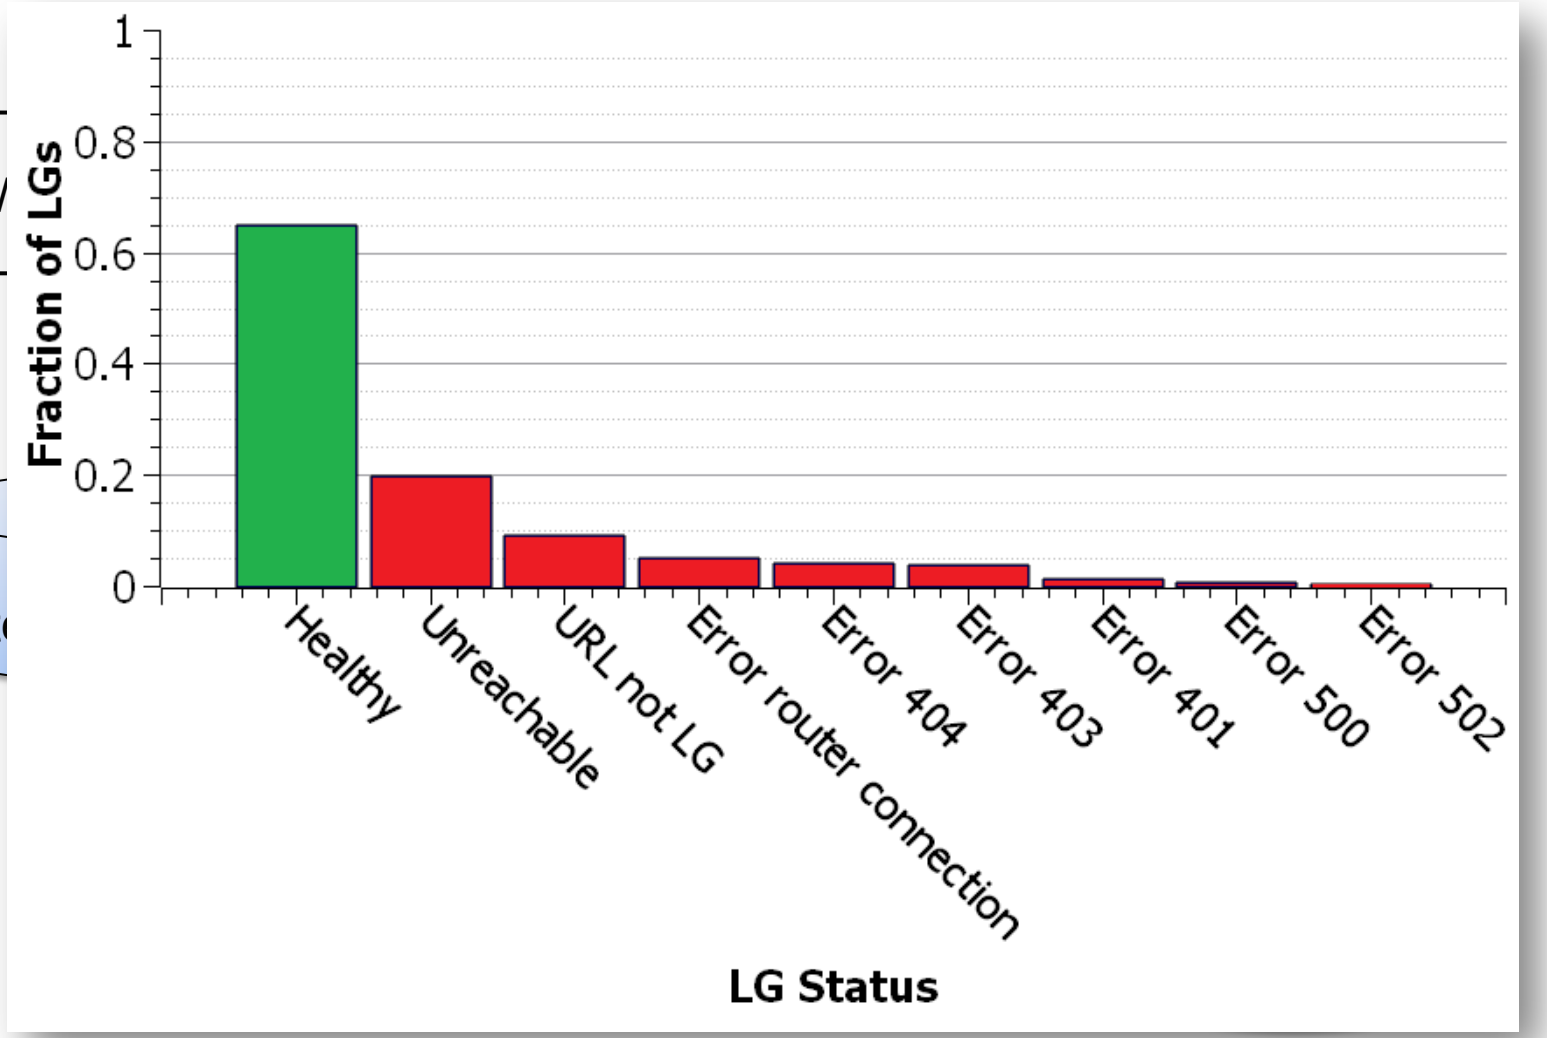

LG Ingestion Workflow

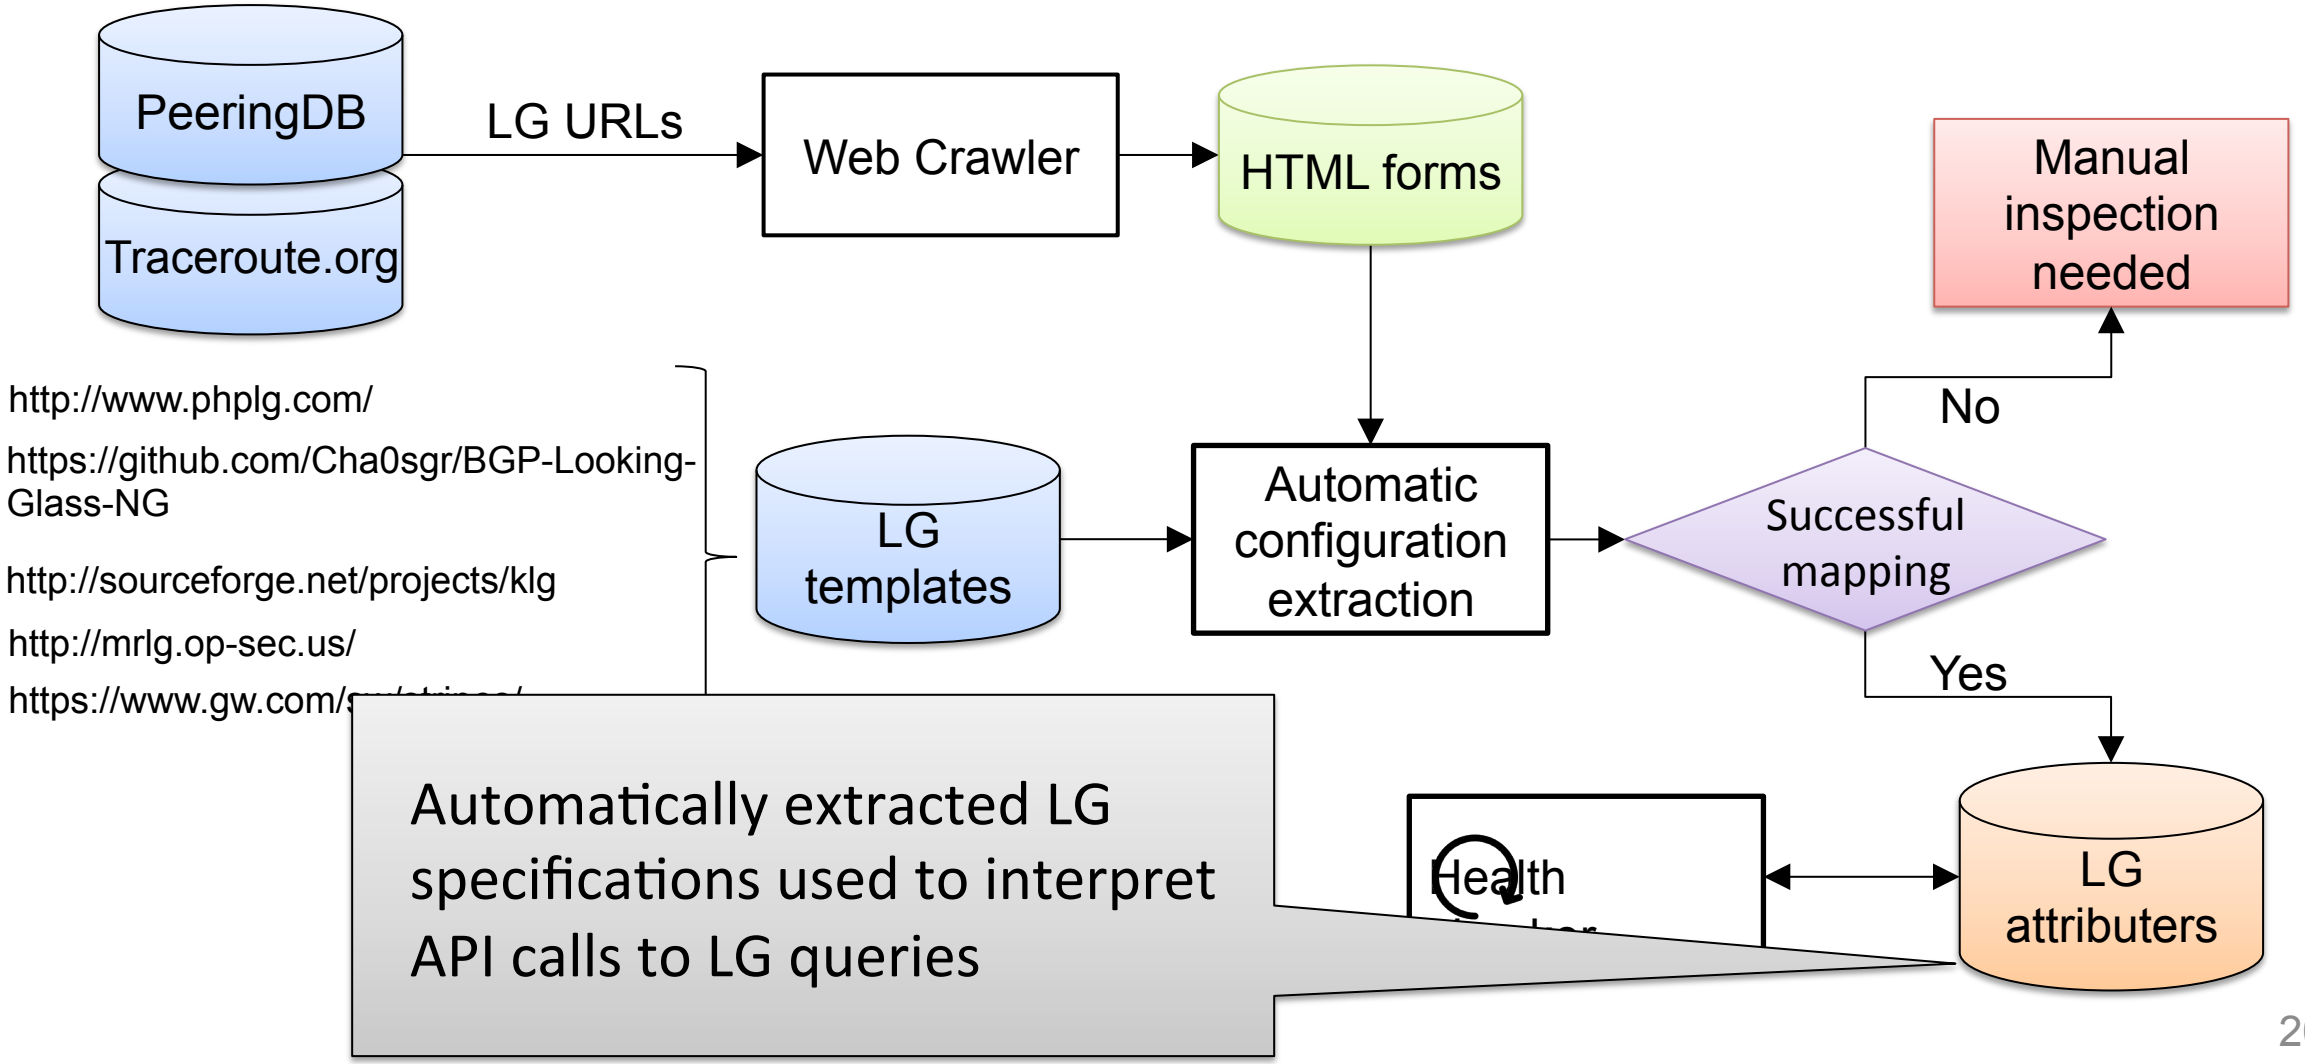

Load Distribution among LGs

```
BGP Routing table entry for 103.22.203.0/24
Paths: (2 available, best #1)

 6453 13335
  AS-path translation: 103.22.203.0/24
  edge4.Frankfurt1 (metric 0)
   Community: Europe Backbone_2 Lclprf_86 Germany Level3_Peer Frankfurt
   6453:3000 6453:3100 6453:3103
   Origin: IGP, metric 0, localpref 86, Used  Valid  Best  IGP  Group-Best
   Originator: edge4.Frankfurt1

 6453 13335
 AS-path translation: 103.22.203.0/24
  edge4.Frankfurt1 (metric 0)
   Community: Europe Backbone_2 Lclprf_86 Germany Level3_Peer Frankfurt
    6453:3000 6453:3100 6453:3103
   Origin: IGP, metric 0, localpref 86, Valid  IGP
   Originator: edge4.Frankfurt1
```

AS3356 (High query load)

```
BGP routing table entry for 103.22.203.0/24, version 126601054
BGP Bestpath: deterministic-med
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
     8
  Refresh Epoch 1
  3356 3356 6453 13335
    188.1.200.77 (metric 1141) from 188.1.200.81 (188.1.200.81)
      Origin IGP, metric 0, localpref 80, valid, internal
      Community: 680:66 3356:2 (Europe) 3356:22 3356:86 3356:501 (Germany)
                 3356:666 (Peer route) 3356:2065 (FRF - Frankfurt) 6453:3000
                 6453:3100 6453:3103
      Originator: 188.1.200.77, Cluster list: 188.1.200.81
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  3356 3356 6453 13335
    188.1.200.77 (metric 1141) from 188.1.200.77 (188.1.200.77)
      Origin IGP, metric 0, localpref 80, valid, internal, best
      Community: 680:66 3356:2 (Europe) 3356:22 3356:86 3356:501 (Germany)
      3356:666 (Peer route) 3356:2065 (FRF - Frankfurt) 6453:3000 6453:3100
      6453:3103
```

AS680 (Low query load)

Load Distribution among Platforms

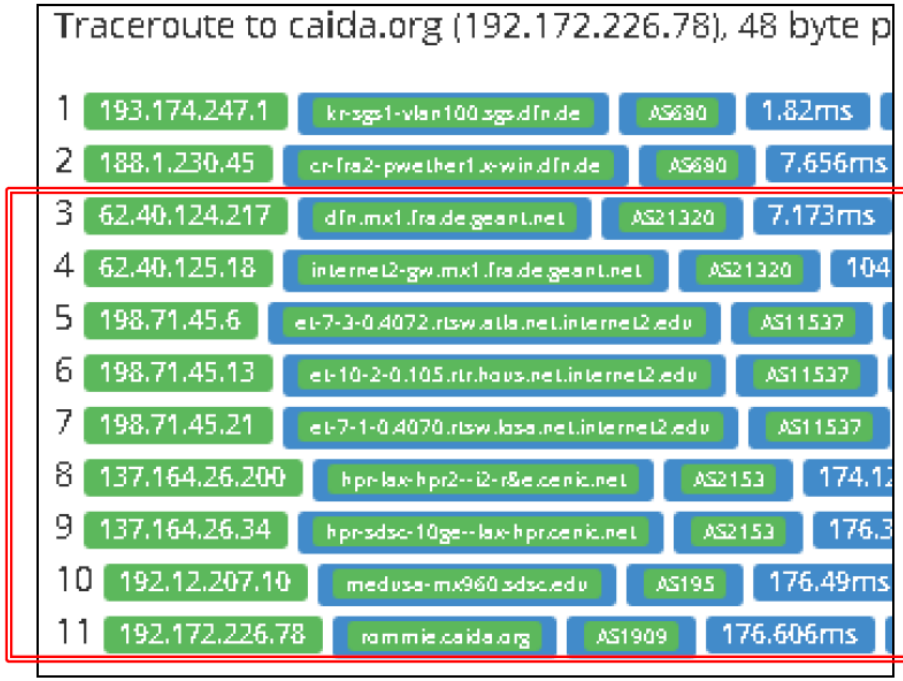

AS680 (RIPE Atlas probe)

```
Type escape sequence to abort.
Tracing the route to ns1.caida.org (192.172.226.78)
VRF info: (vrf in name/id, vrf out name/id)
  1 xr-fzk1-pc2.x-win.dfn.de (188.1.145.81) [MPLS: Label 1274
  2 cr-fra2-be9.x-win.dfn.de (188.1.144.121) 4 msec 8 msec 4 r
  3 dfn.mx1.fra.de.geant.net (62.40.124.217) [AS 20965] 4 msec
  4 internet2-gw.mx1.fra.de.geant.net (62.40.125.18) [AS 20965
  5 et-7-3-0.4072.rtsw.atla.net.internet2.edu (198.71.45.6) [A
  6 et-10-2-0.105.rtr.hous.net.internet2.edu (198.71.45.13) [A
  7 et-7-1-0.4070.rtsw.losa.net.internet2.edu (198.71.45.21) |
  8 hpr-lax-hpr2--i2-r&e.cenic.net (137.164.26.200) [AS 2153]
  9 hpr-sdsc-10ge--lax-hpr.cenic.net (137.164.26.34) [AS 2153]
 10 medusa-mx960.sdsc.edu (192.12.207.10) [AS 195] 180 msec 1
 11 ns1.caida.org (192.172.226.78) [AS 1909] 188 msec 176 msec
```

AS680 (Looking Glass)

Case study: ARTEMIS

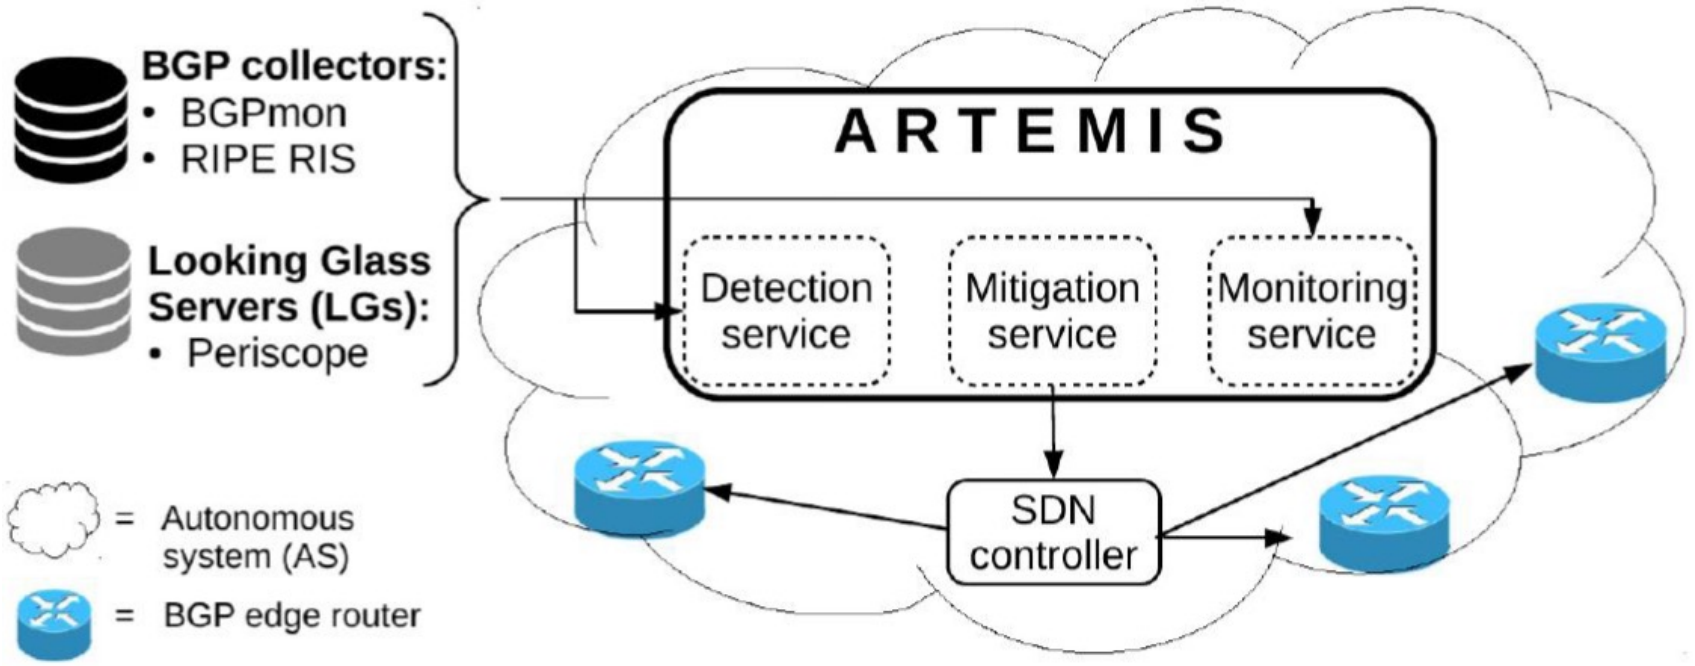

G. Chaviaras et al. ARTEMIS: Real-Time Detection and Automatic Mitigation for BGP Prefix Hijacking. *ACM SIGCOMM 2016*