PCH
Packet Clearing House

Visualizing a global DNS network with open-source tools

Ashley Jones
Packet Clearing House

APRICOT 2018
Kathmandu, Nepal

pch.net/OSSViz

PCH
Packet Clearing House

Who are we?

The international, non-profit organization responsible for providing operational support and security to critical Internet infrastructure, including Internet exchange points and the core of the domain name system.

PCH
Packet Clearing House

Problem Statement:

"How can PCH improve statistics for DNS services we host?"

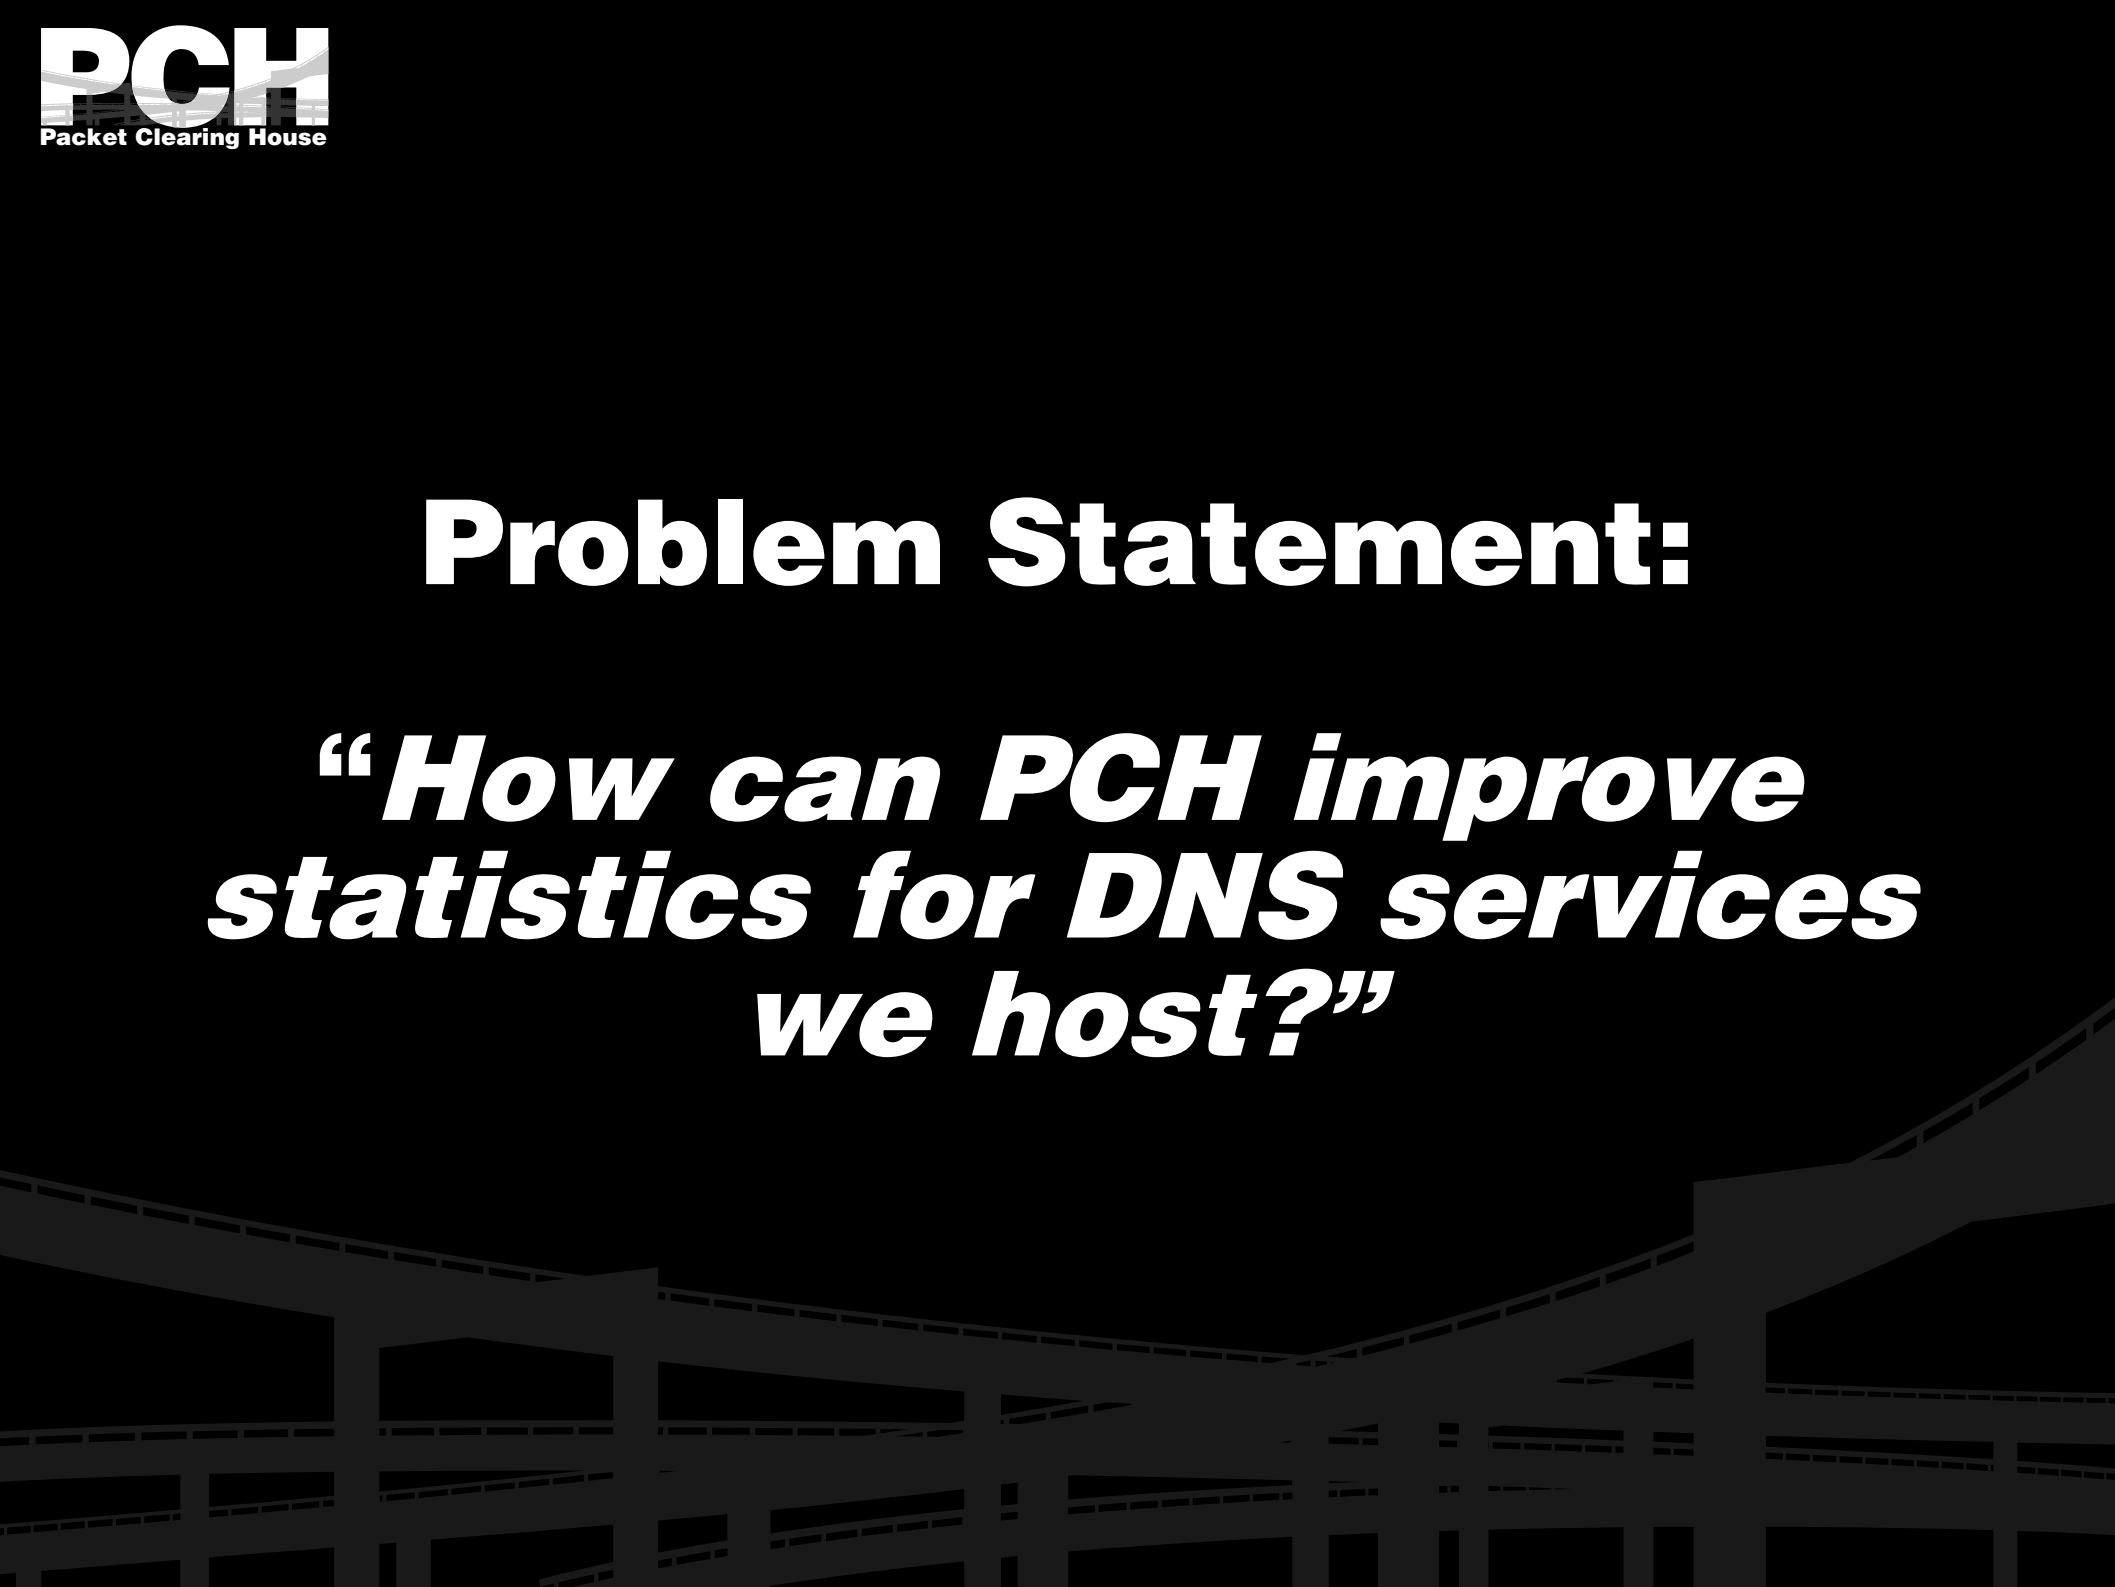

PCH
Packet Clearing House

Stats are from fake ".aj" TLD

City	Queries	Queries per Second	Query Share	Responses	Resp per Second	Resp Share
Kathmandu	2088734	24	28%	2088941	24	28%
Ashburn	1053708	12	14%	1056436	12	14%
Frankfurt	871876	10	12%	871733	10	12%
New York	647887	7.5	8.6%	647716	7.5	8.6%
London	370744	4.3	4.9%	370703	4.3	4.9%
Seattle	355408	4.1	4.7%	355330	4.1	4.7%
Singapore	251966	2.9	3.3%	252870	2.9	3.3%
Miami	203937	2.4	2.7%	203838	2.4	2.7%
Chicago	203097	2.4	2.7%	207343	2.4	2.7%
Palo Alto	146469	1.7	1.9%	146328	1.7	1.9%
Los Angeles	140605	1.6	1.9%	140596	1.6	1.9%
Cape Town	138703	1.6	1.8%	138459	1.6	1.8%
Dublin	129810	1.5	1.7%	129804	1.5	1.7%
Atlanta	96326	1.1	1.3%	96327	1.1	1.3%
Tokyo	81864	0.95	1.1%	81840	0.95	1.1%
Hong Kong	80837	0.94	1.1%	73664	0.85	0.97%
Sydney	66803	0.77	0.88%	66793	0.77	0.88%
Jakarta	66437	0.77	0.88%	66262	0.77	0.87%

DNS Stats - .aj - Current vs Baseline By City

Download

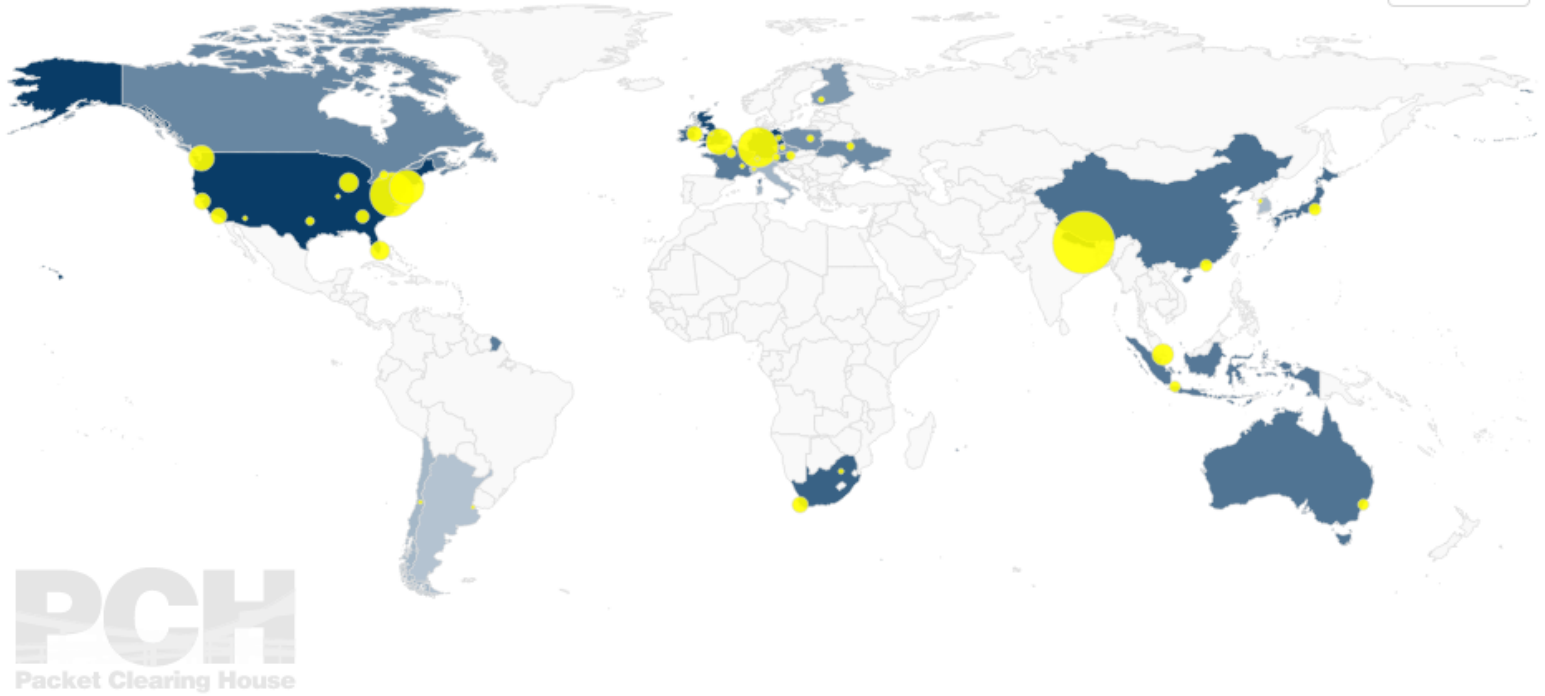

Each marker represents a city. Color of country represents the sum of queries in that country. Source: **pch.net**

Filters

+ New filter

		QPS		TCP %		IPv6 %	
City ▲	Country	Current	Baseline	Current	Baseline	Current	Baseline
Ashburn	US	1106 (5%)	2762 (31%)	5%	5%	4%	3%
Atlanta	US	8424 (2%)	8749 (31%)	5%	2%	4%	3%
Berlin	DE	4654 (4%)	7982 (01%)	2%	2%	2%	4%

PCH
Packet Clearing House

This space intentionally blank because PCH has no graphs right now ;)

PCH
Packet Clearing House

DNS Stats - .aj - Last 24 hours

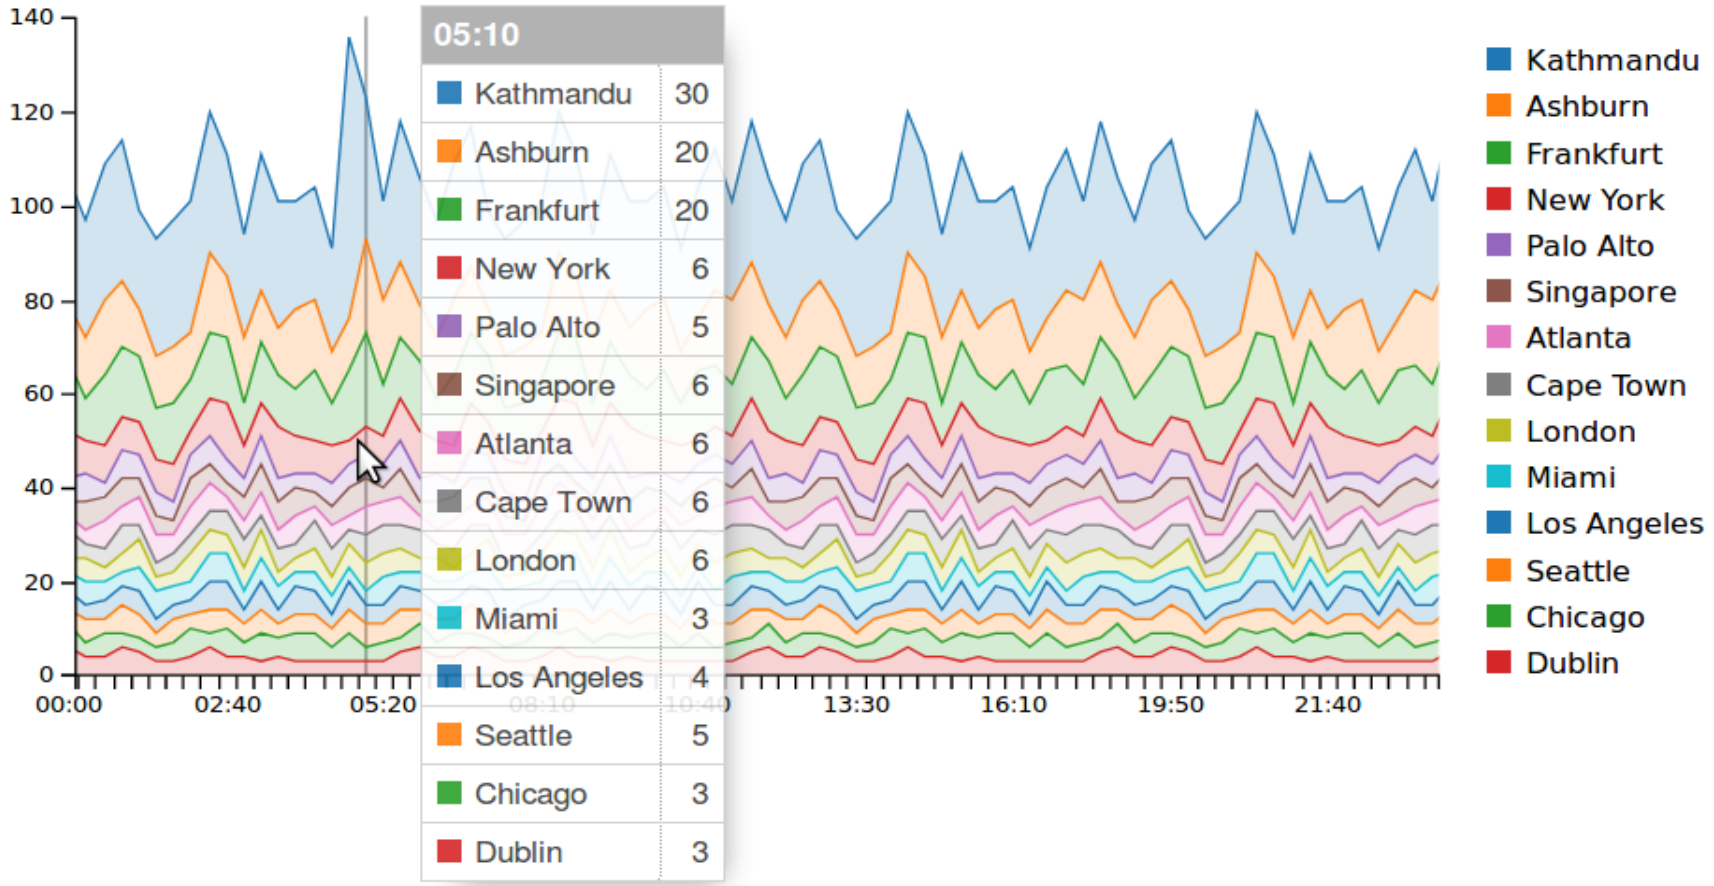

PCH
Packet Clearing House

PCH ccTLD DNS Service

- 130+ PoPs
- Anycast

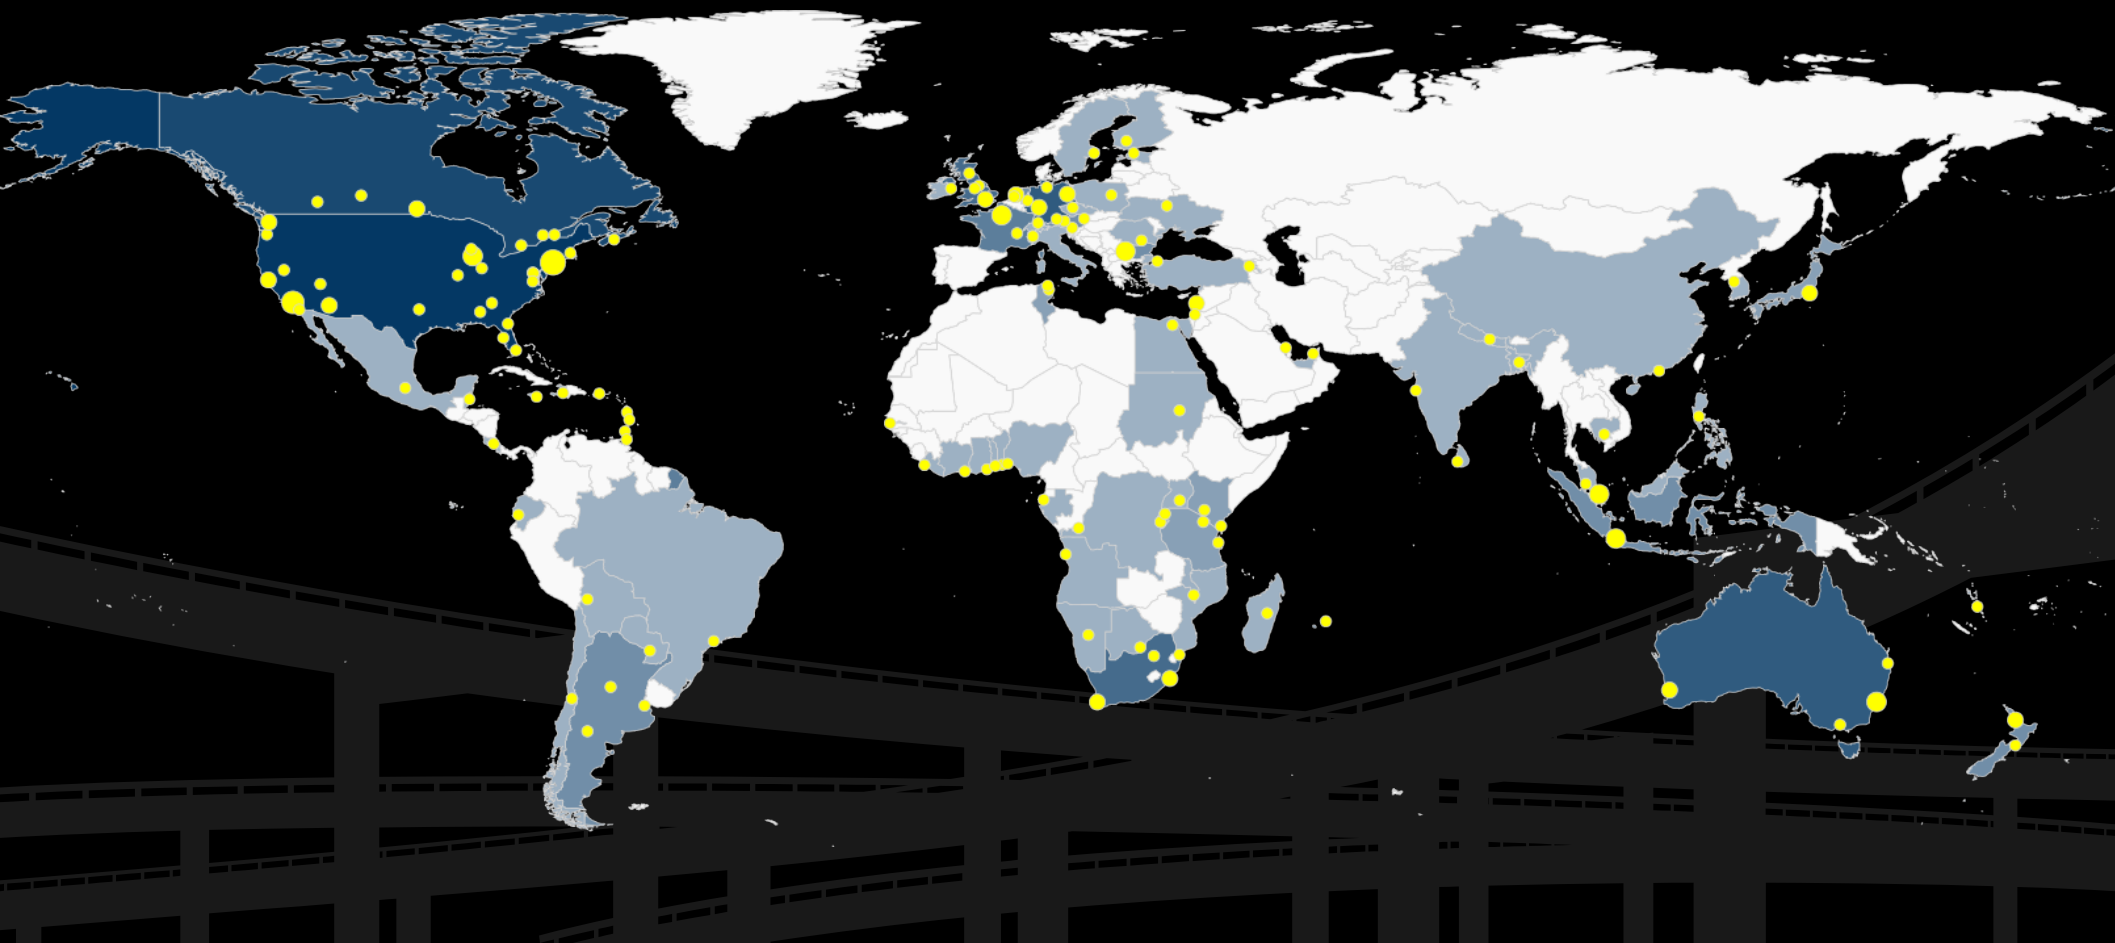

PCH
Packet Clearing House

PCH ccTLD DNS Service

- 109 ccTLDs served

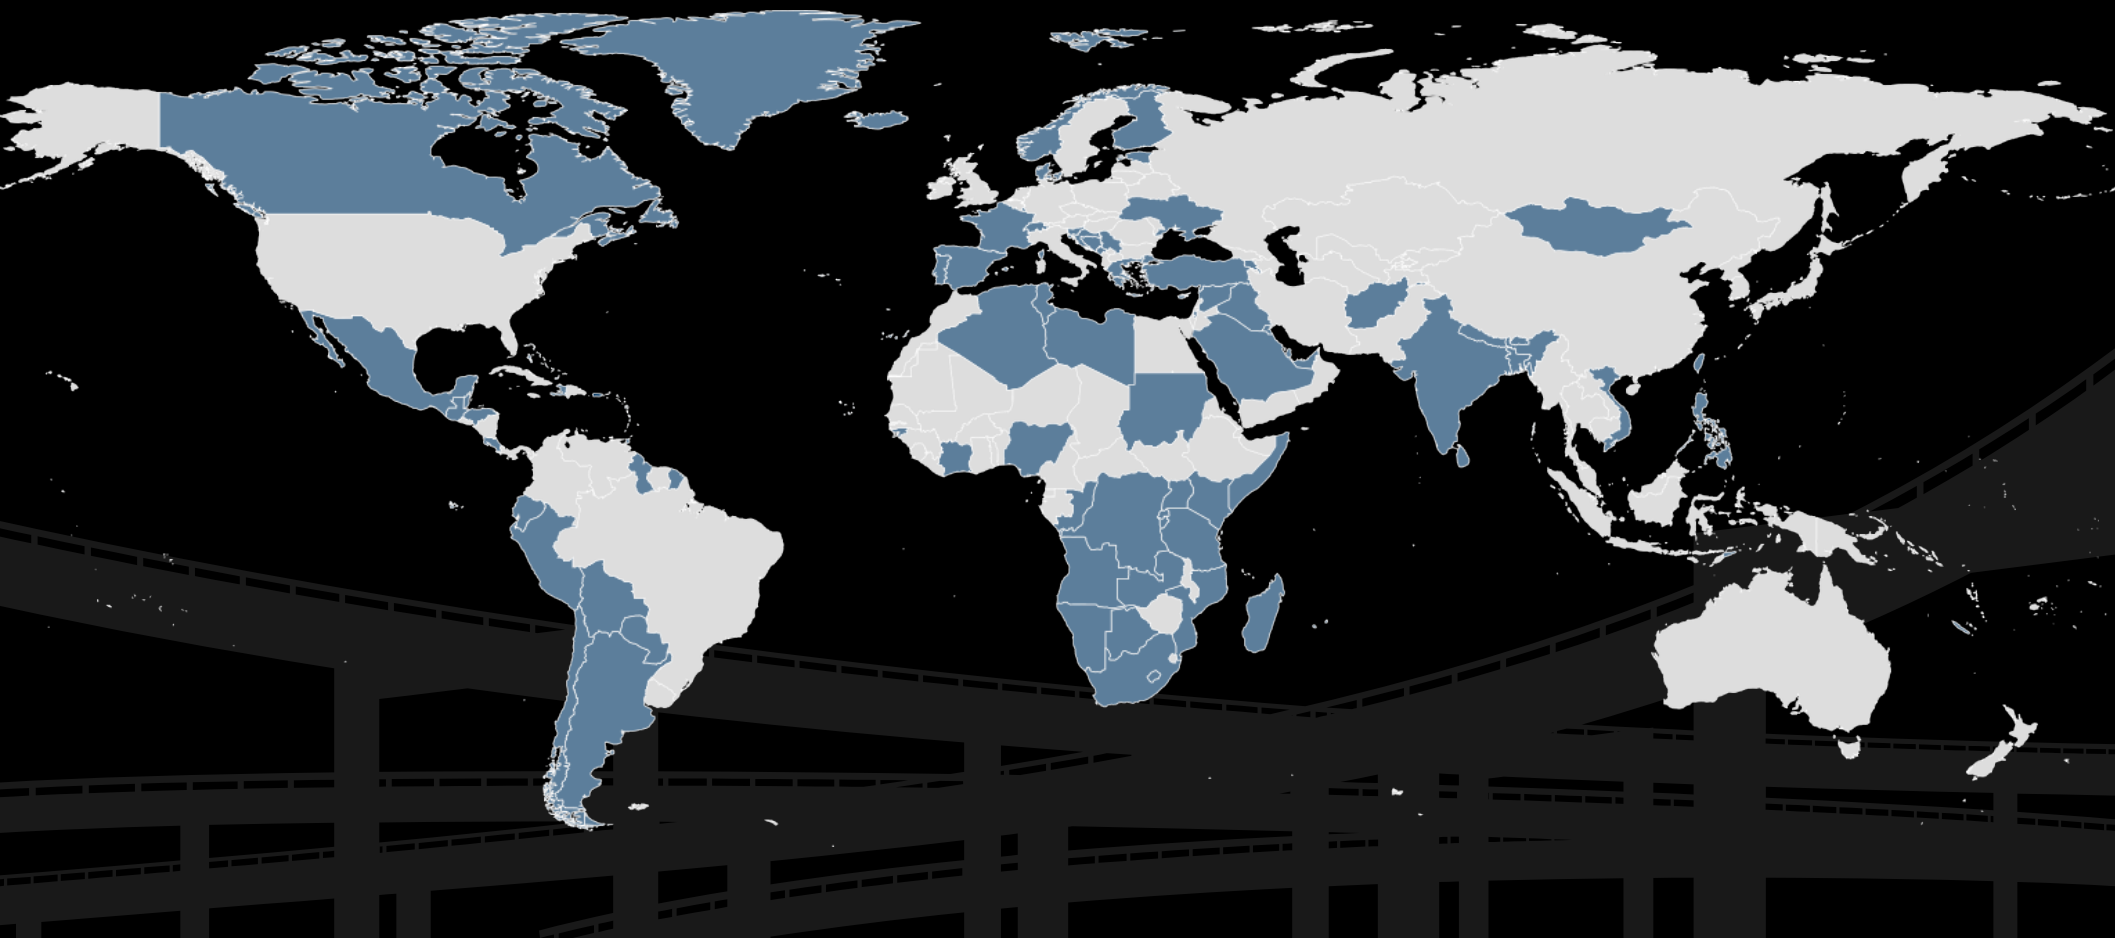

PCH
Packet Clearing House

DNS Node Technology

- VMware ESXi on Cisco hardware
- Multiple VMs running BIND & NSD
- Public Anycast IPs

PCH
Packet Clearing House

DNS Network Topology

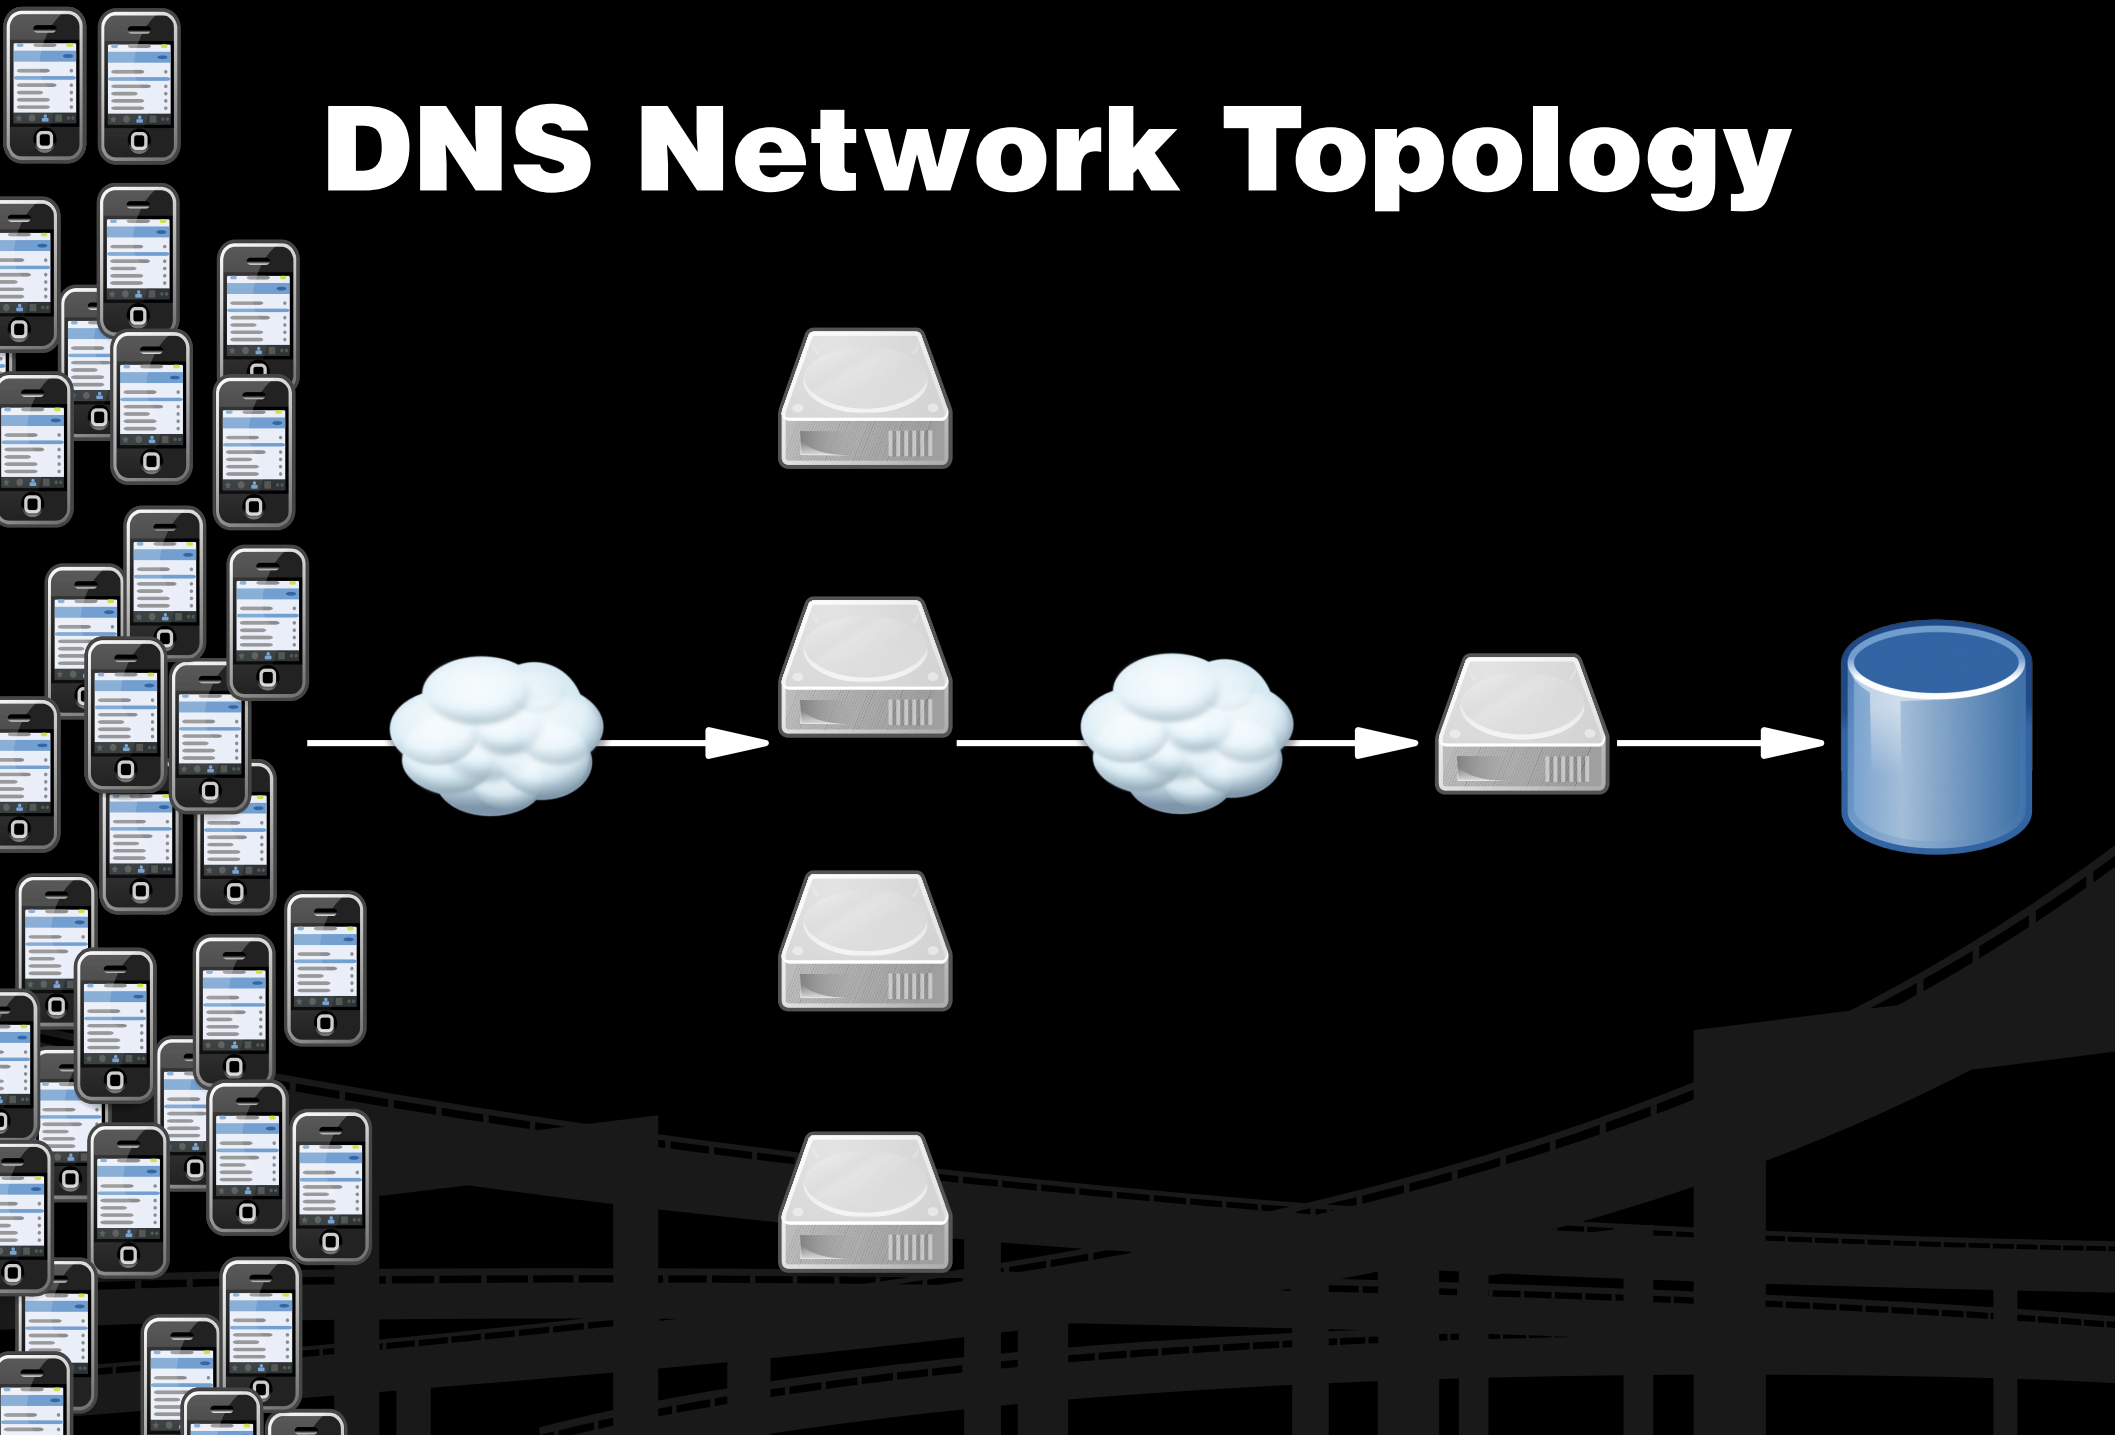

PCH
Packet Clearing House

DNS Network Topology

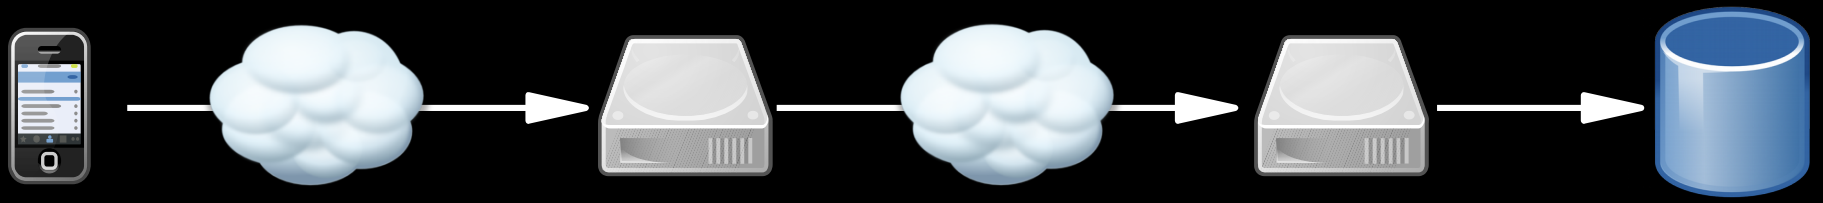

PCH
Packet Clearing House

Existing Process

1. Distill pcap at the node with java
2. Push distilled pcaps to master
3. 2nd java app cooks pcaps into a MySQL row
4. Avoid slow DB inserts via table/day with ~20M rows of zones/city/10min (No full zones queried stored, only customer ccTLD)

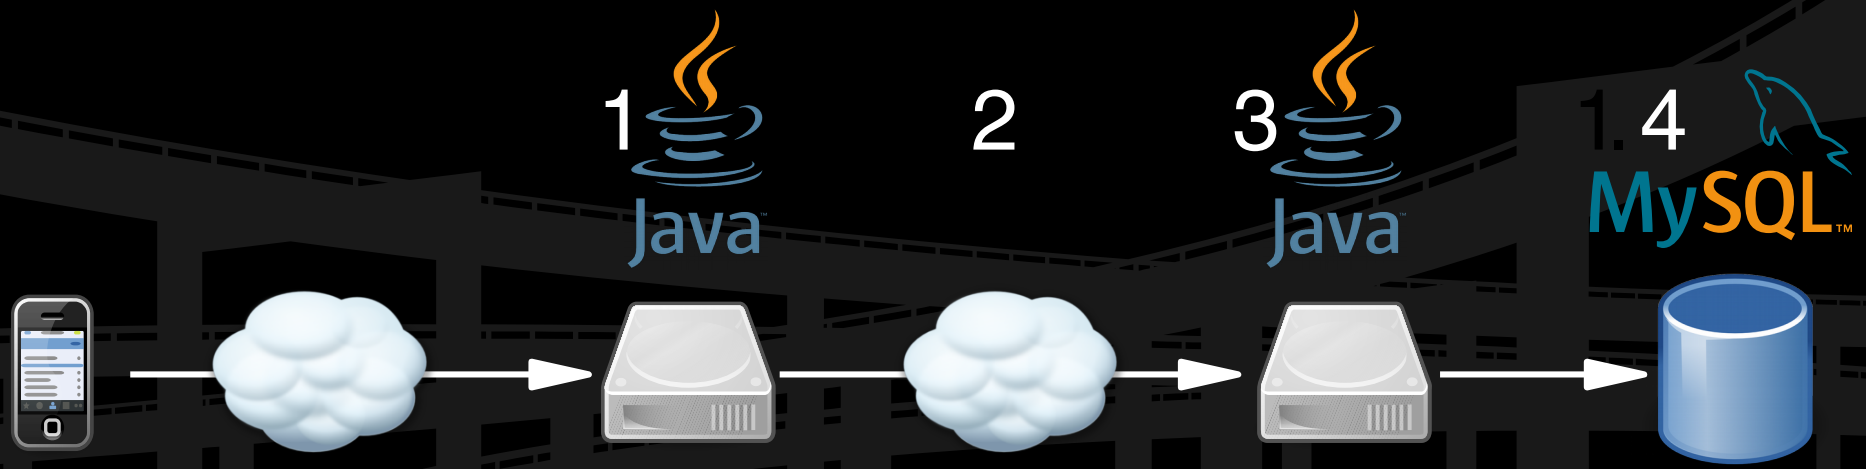

PCH
Packet Clearing House

Problems with Existing Stats

- Slow to view/download (30+ seconds)
- 10 minute granularity
- Only tabular
- Static HTML

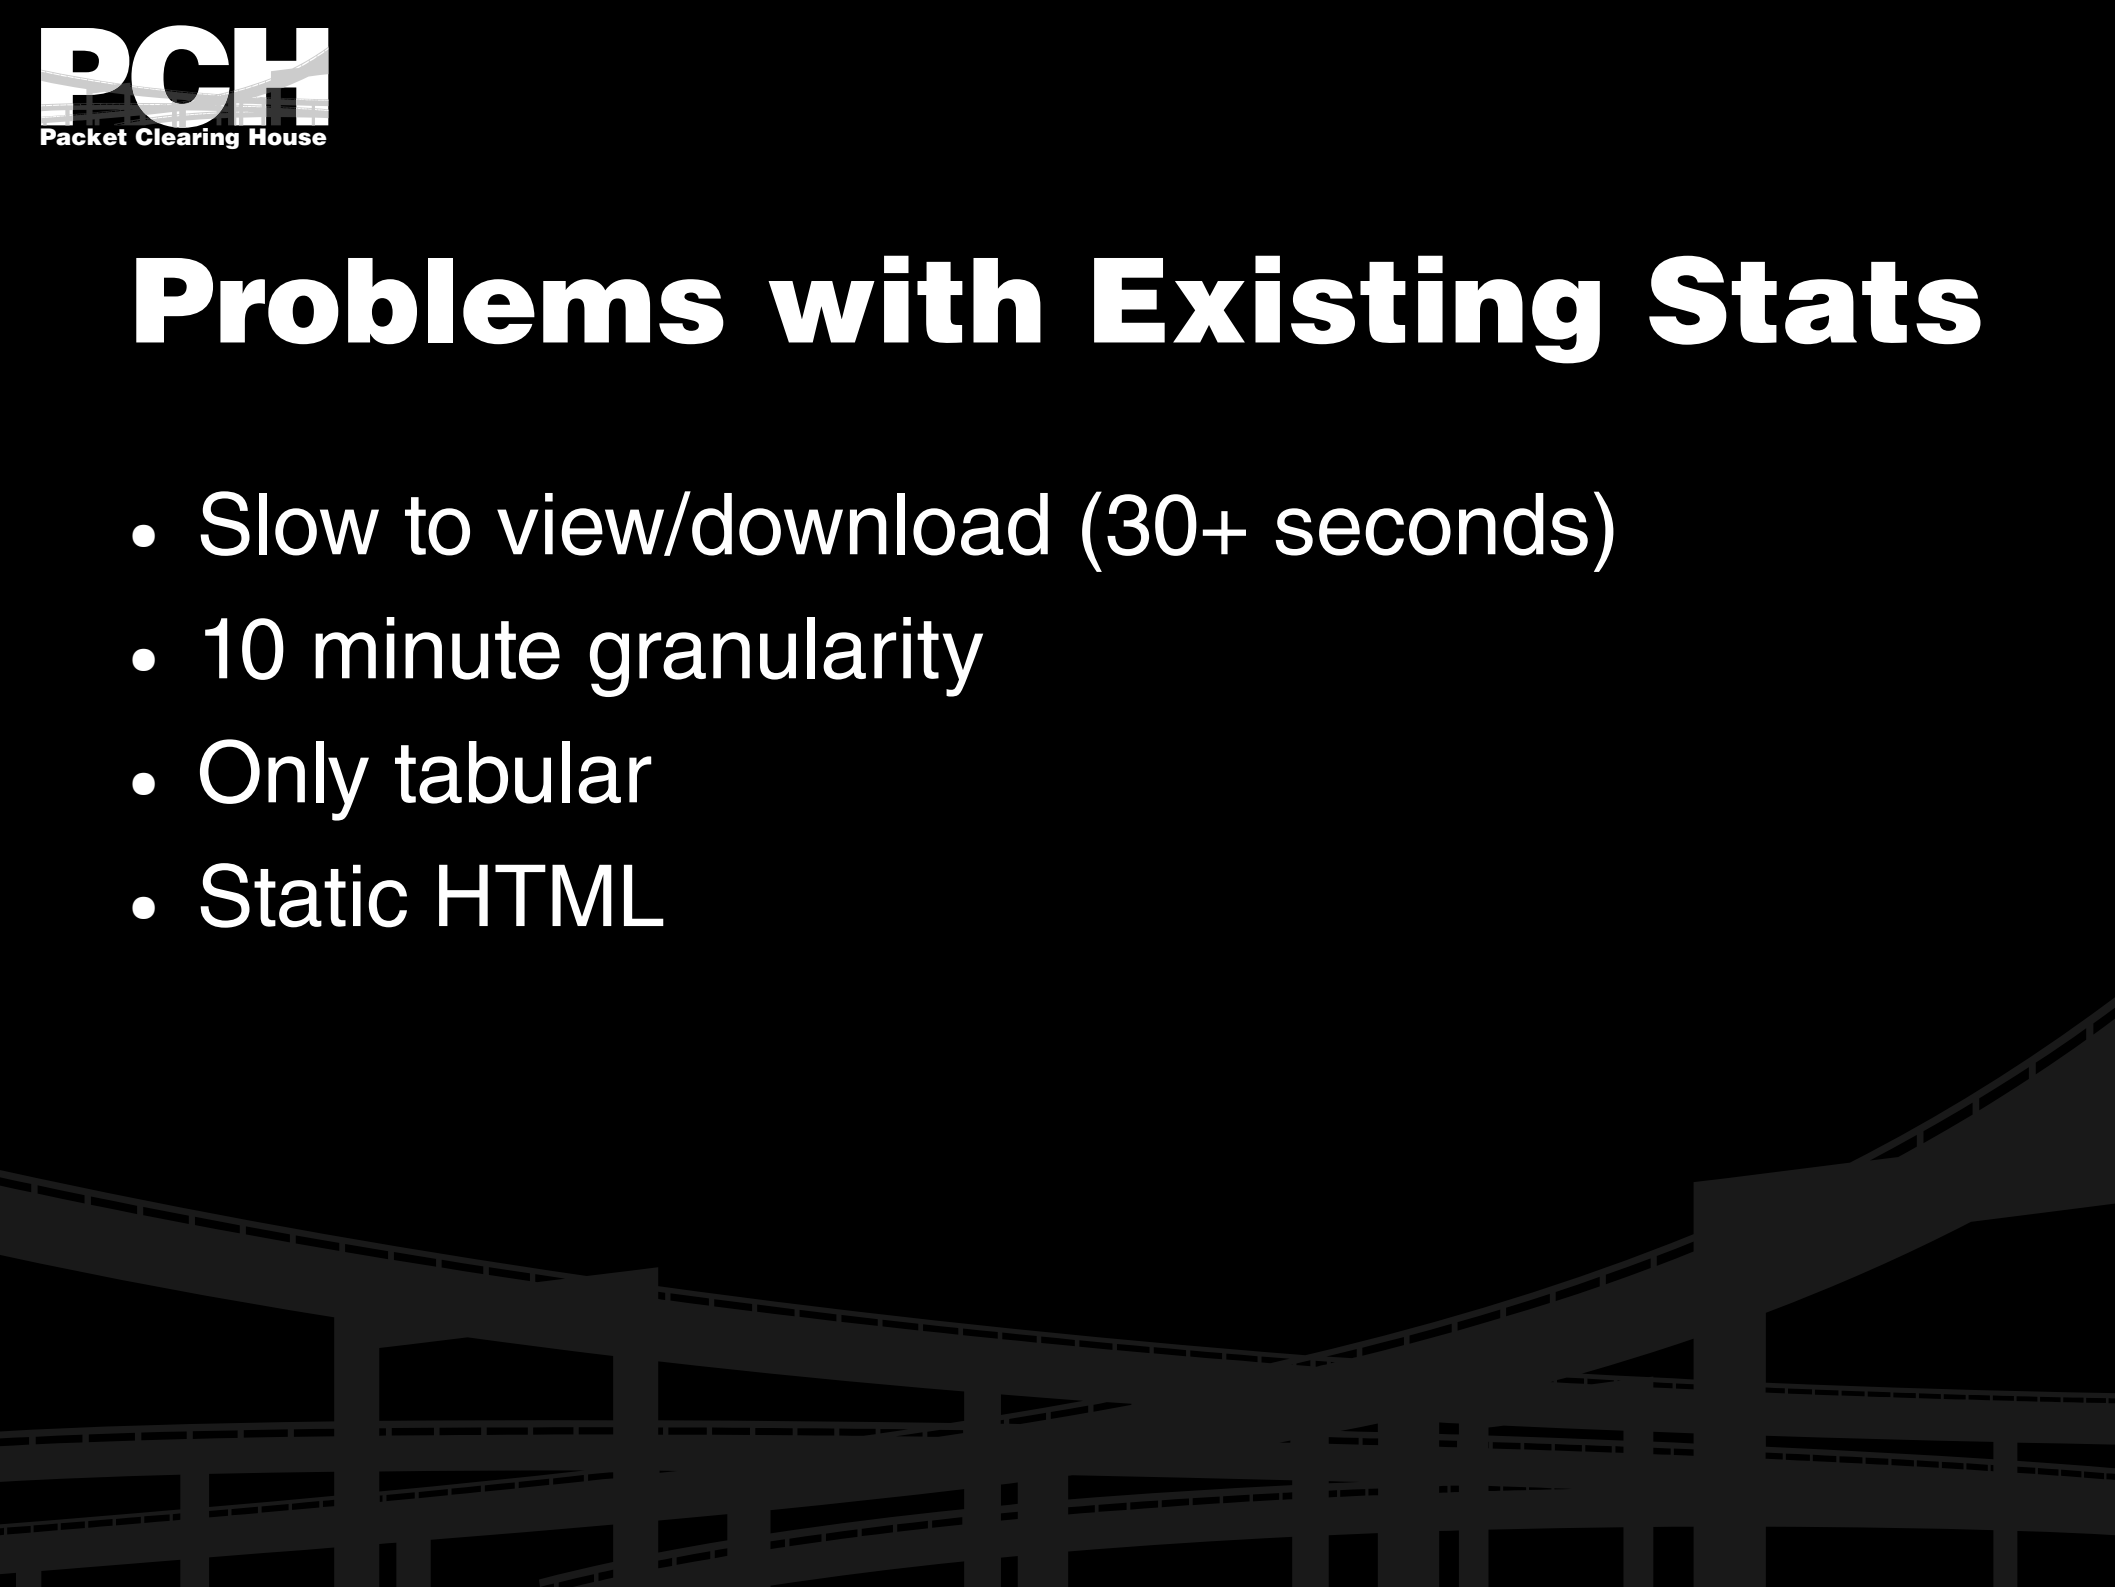

PCH
Packet Clearing House

Problem Statement

"How can PCH improve statistics for DNS services we host?"

- Customer facing:
  - Graphs & maps
  - Near real time
  - Load instantly
  - Interactive
  - Downloadable/exportable

PCH
Packet Clearing House

Problem Statement

"How can PCH improve statistics for DNS services we host?"

- Server side:
  - Speed up parsing of pcap
  - Speed up DB reads/writes
  - Ease aggregating data to day → week → month → year

PCH
Packet Clearing House

Solution?

PCH
Packet Clearing House

Open Source!

1. Distill pcap at the node with java
2. Push distilled pcaps to master
3. DNSAuth – PCH Go App for cooking pcaps to TSDB faster
4. InfluxDB – TSDB to archive data

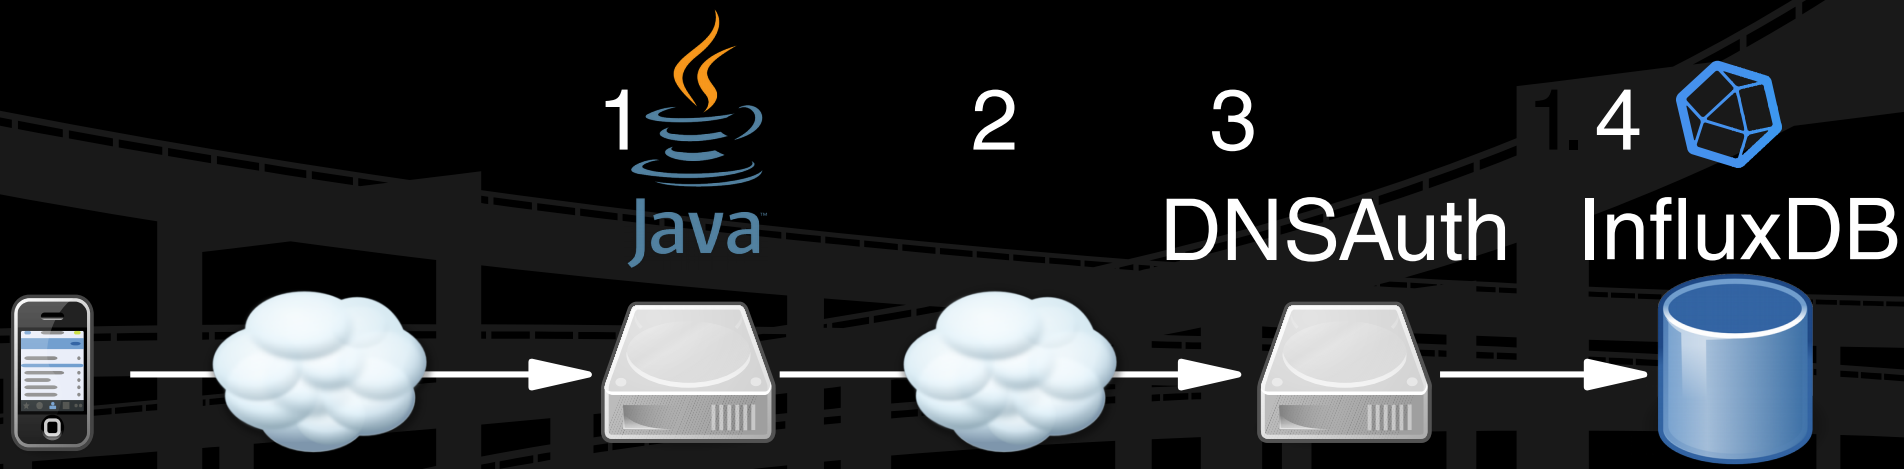

PCH
Packet Clearing House

DNSAuth

- Originally written by PCH for Quad9
- Parse DNS stats as fast as possible
- Written in Go

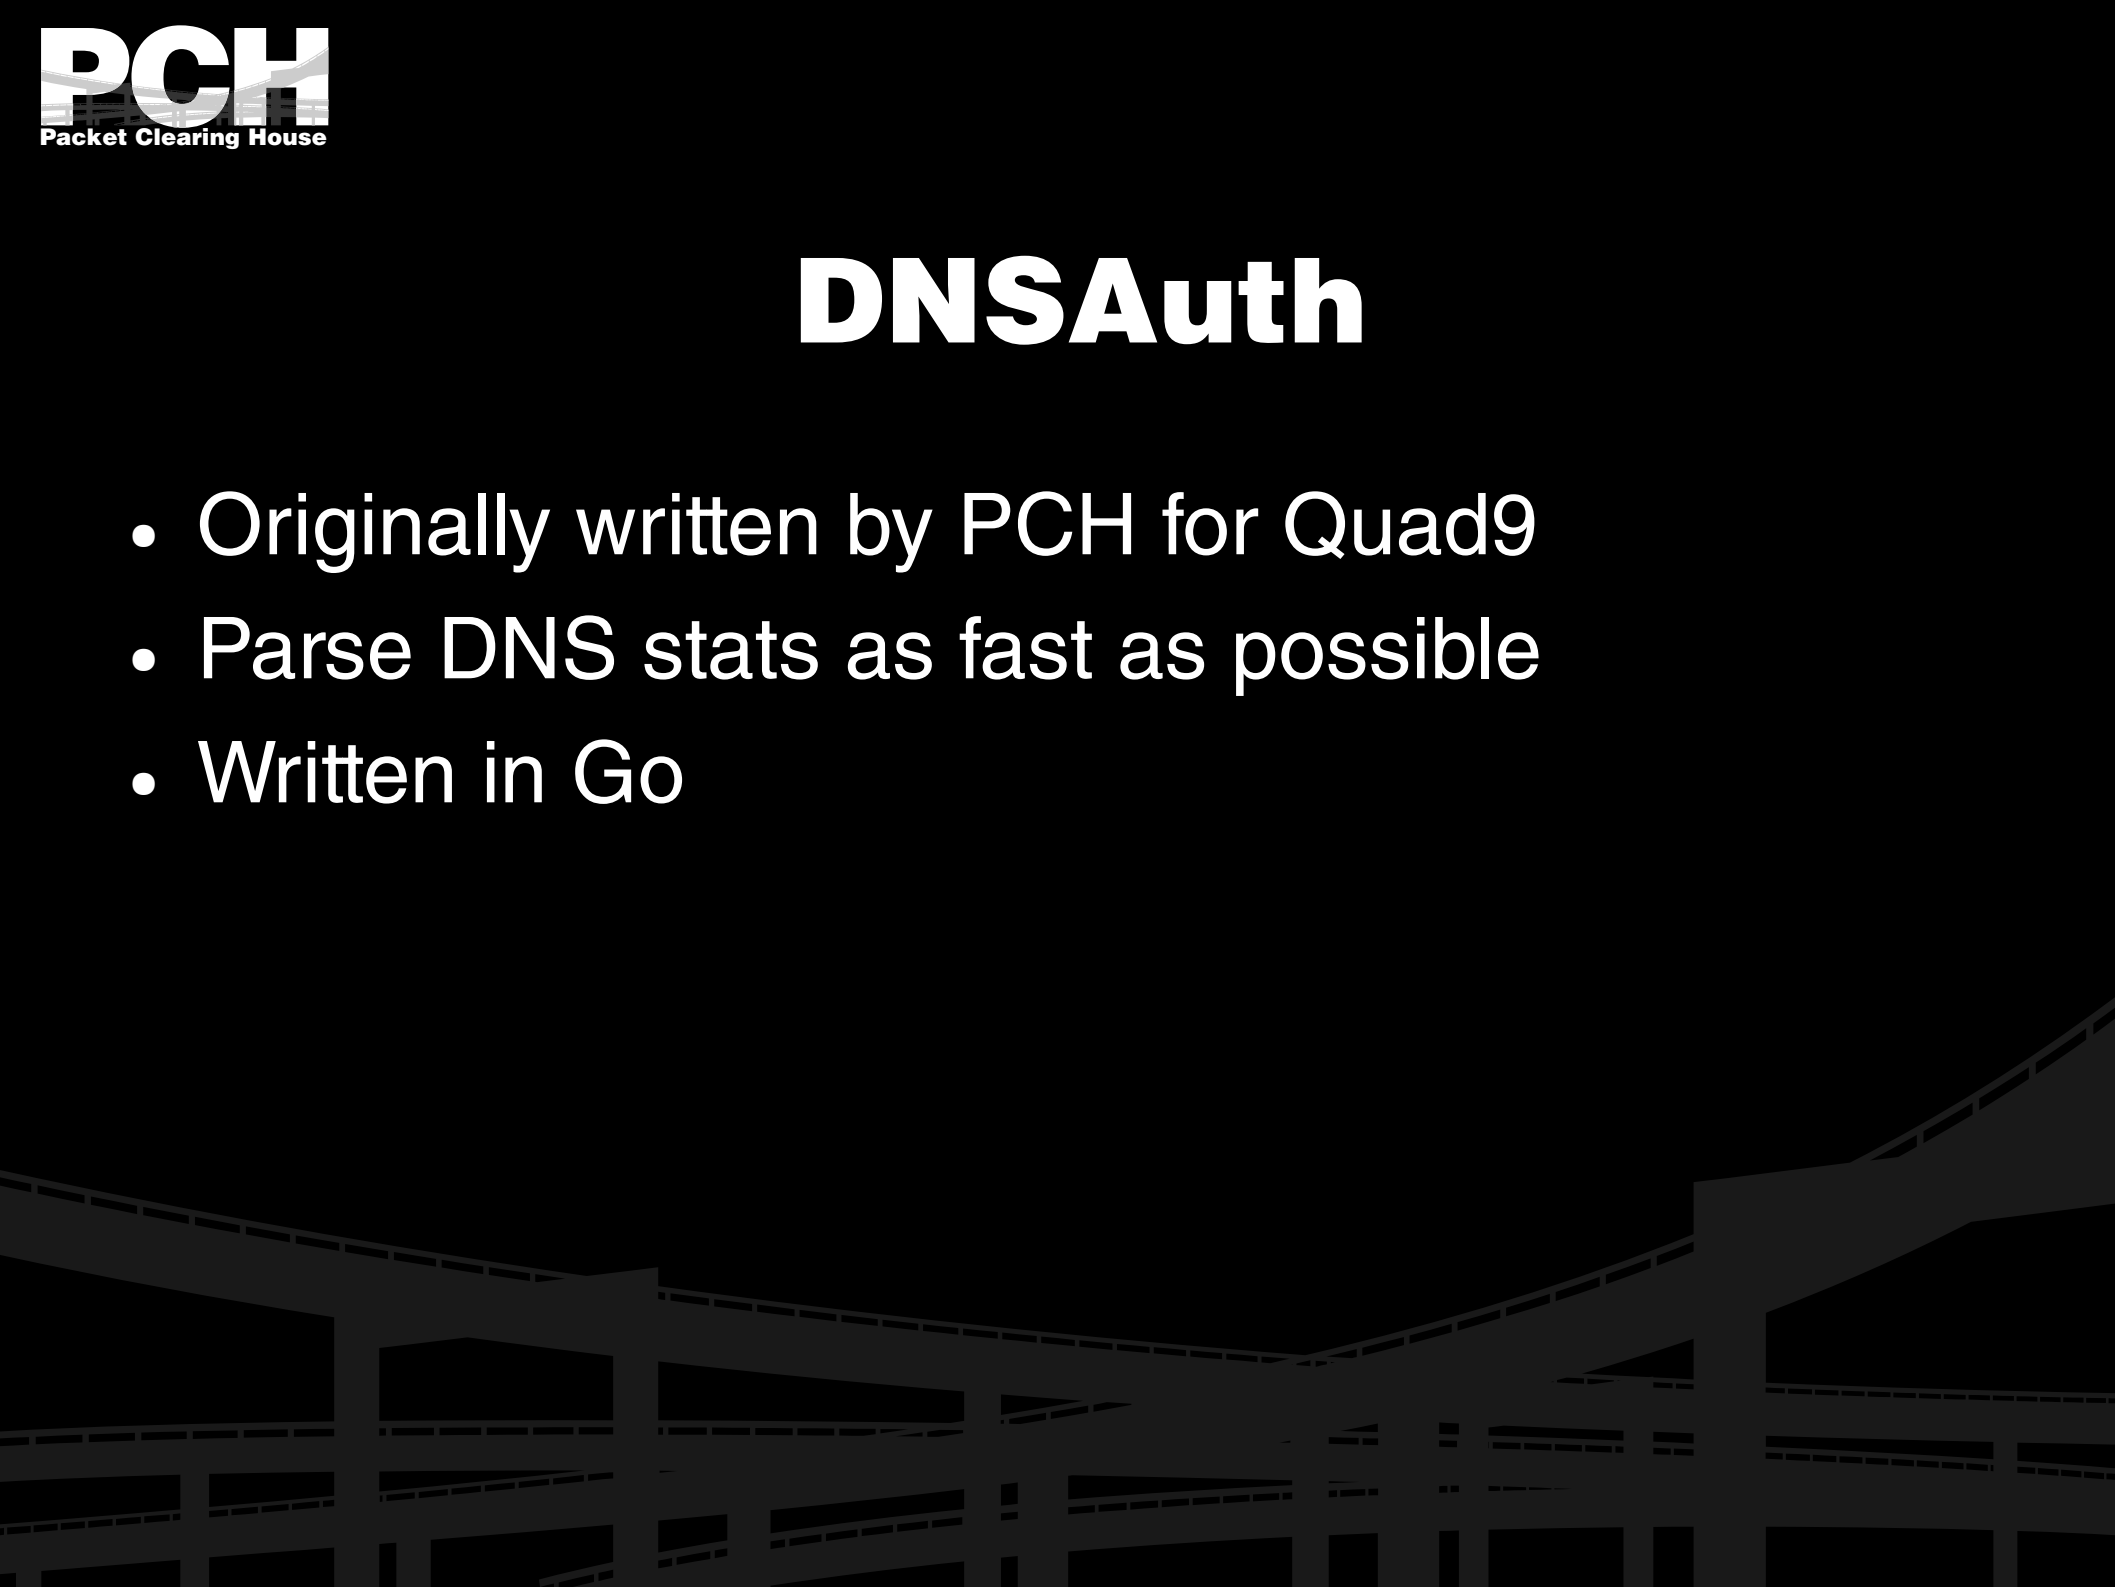

PCH
Packet Clearing House

InfluxDB

- Existing, Mature TSDB
- SDKs for PHP et al.
- FAST!
- Auto aggregating archives to reduce storage demands

PCH
Packet Clearing House

c3

- A D3-based reusable chart library
- Easy to integrate with our PHP app
- Easy to feed real time data from InfluxDB
- Mobile friendly rendering

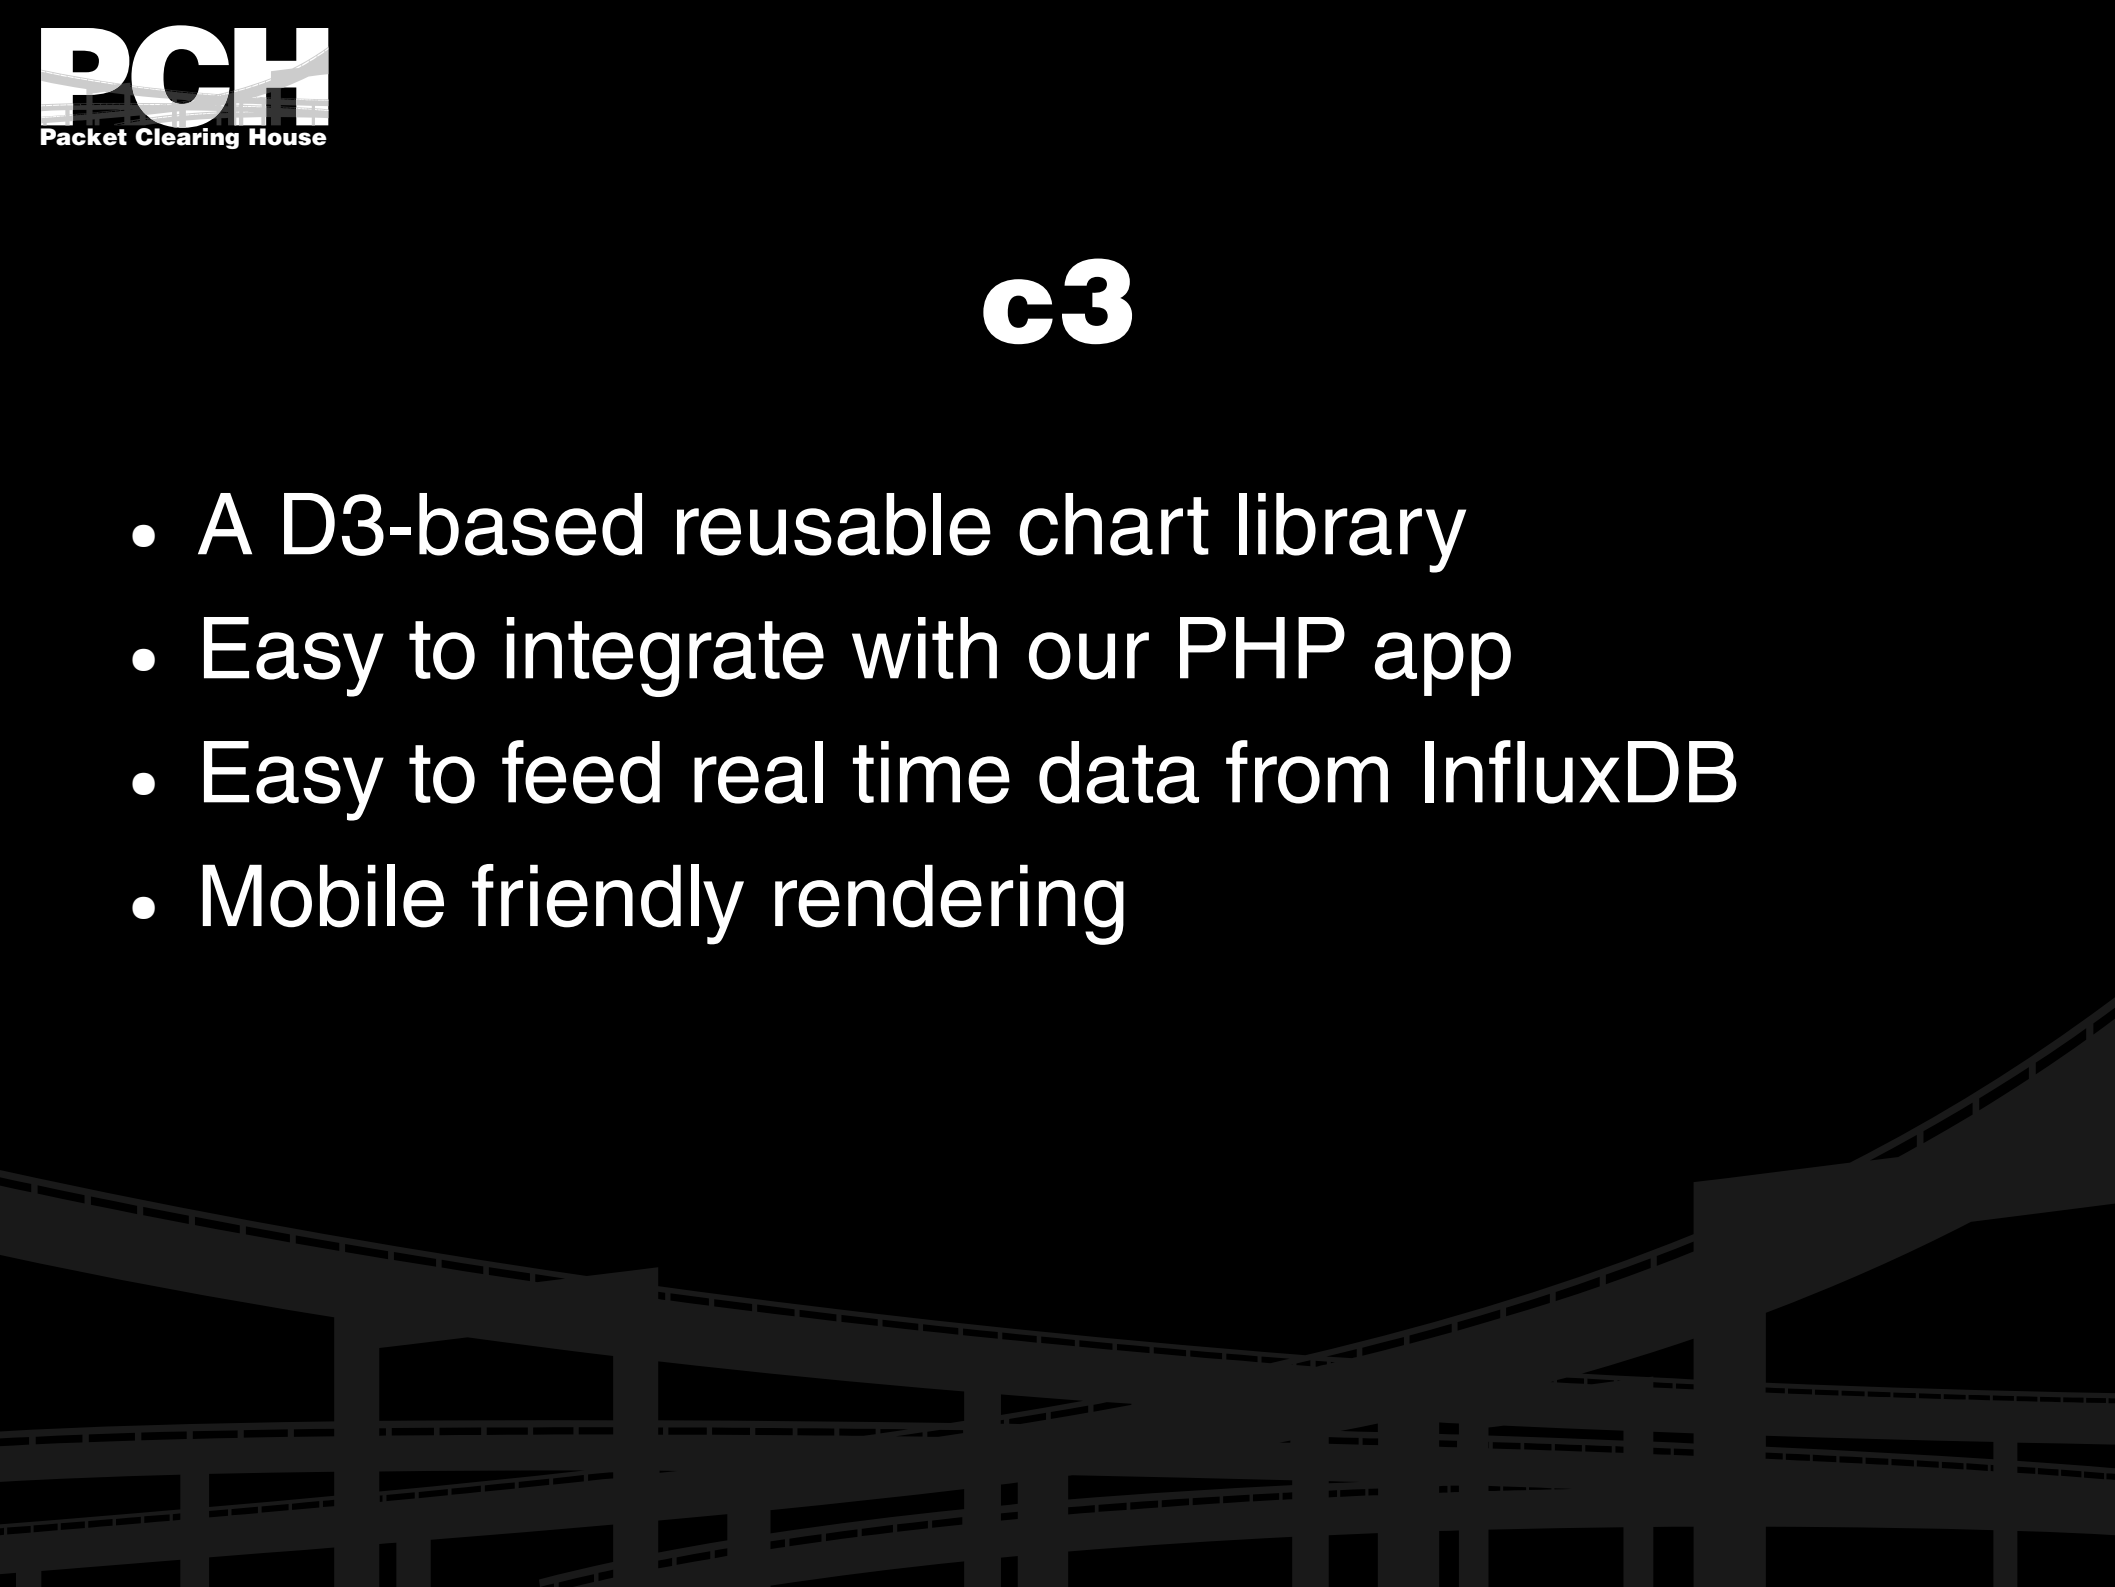

PCH
Packet Clearing House

MapTable

- Written by PCH for IXP Directory
- Easy maps from CSV or JSON
- Customizable including virtual columns & calculated values
- Zoomable, filterable & downloadable.
- Real time filters update map and table in .

PCH
Packet Clearing House

Demo!

(or stills in case it fails ;)

PCH
Packet Clearing House

Next Steps

- Query Graphs (A, NS, CNAME etc.)
- Response Graphs (NOERROR, NXDOMAIN, & SERVFAIL)
- Easier filtering by domains on ccTLD (eg gov.aj vs com.aj)

PCH
Packet Clearing House

Thanks!

Ashley Jones
mrjones@pch.net

pch.net/OSSViz

PCH
Packet Clearing House

URLs

- DNSAuth: github.com/Packet-Clearing-House/DNSAuth
- InfluxDB: github.com/influxdata/InfluxDB
- C3: github.com/c3js/c3
- MapTable: github.com/Packet-Clearing-House/MapTable
- Faux-Logs: https://github.com/Packet-Clearing-House/Faux-Logs

PCH
Packet Clearing House

DNSAuth Distilled pcap

```
R 10.0.2.1 10.0.1.2 0 0 15 foo.com. 582 0
```

R – Query or response
10.0.2.1 – Client IP
10.0.1.2 – Server IP
0 – 0 = UDP, 1 = TCP
0 – Operation code 0=Query, 4=Notify, 5=Update, etc.
15 – Query Type; 1=A, 2=NS, 5=CNAME etc
foo.com. – zone being queried
582 – packet size in bytes
0 – Response 0=NOERROR, 3=NXDOMAIN, 2=SERVFAIL, etc.

PCH
Packet Clearing House

Screenshots in case of demo fail

PCH
Packet Clearing House

Logged in as **mrjones**

PCH
Packet Clearing House

Admin | About | Services | Tools | Resources

Home > Account > DNS Stats Graph

DNS Stats - .aj - Last 24 hours

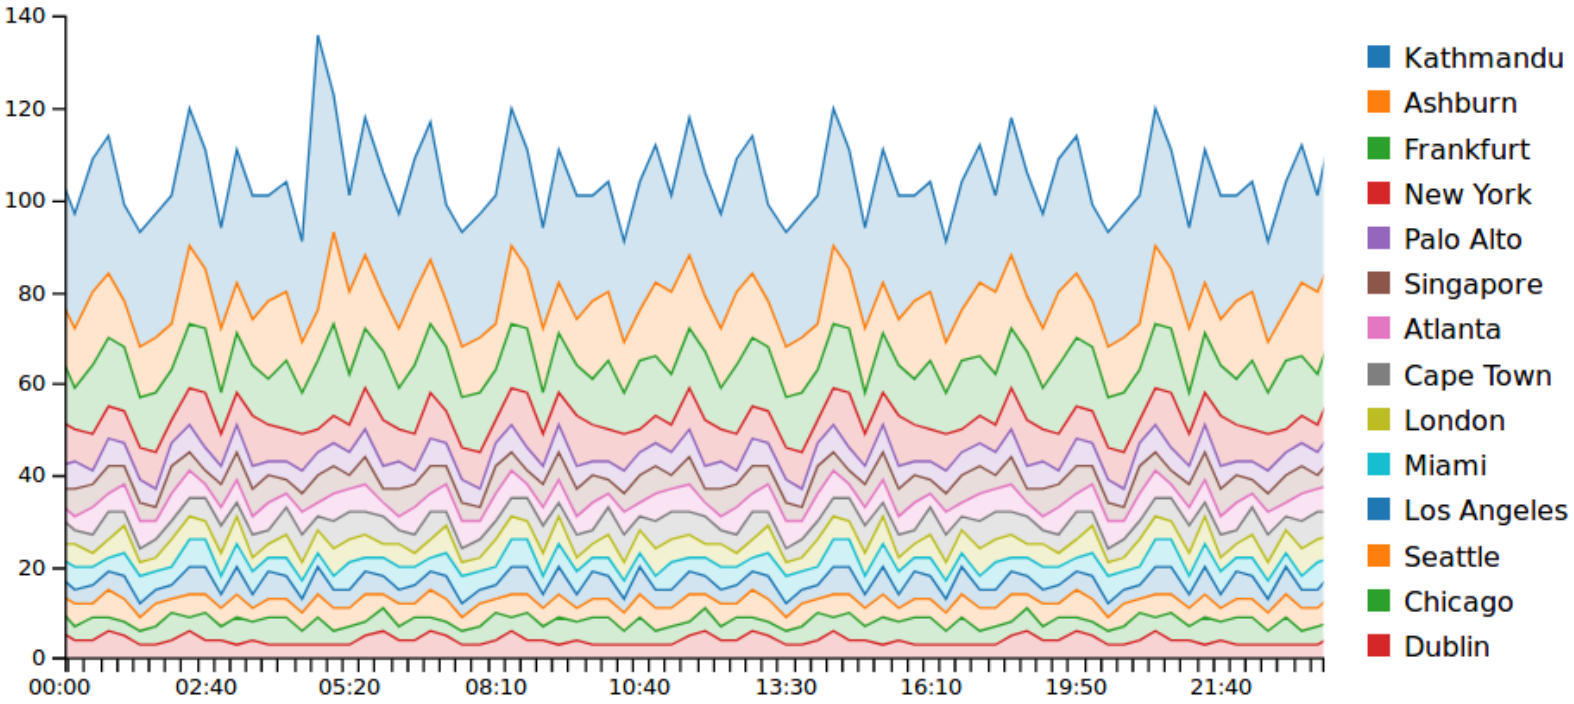

PCH
Packet Clearing House

Logged in as **mrjones**

PCH
Packet Clearing House

Admin | About | Services | Tools | Resources

Home > Account > DNS Stats Graph

DNS Stats - .aj - Last 24 hours

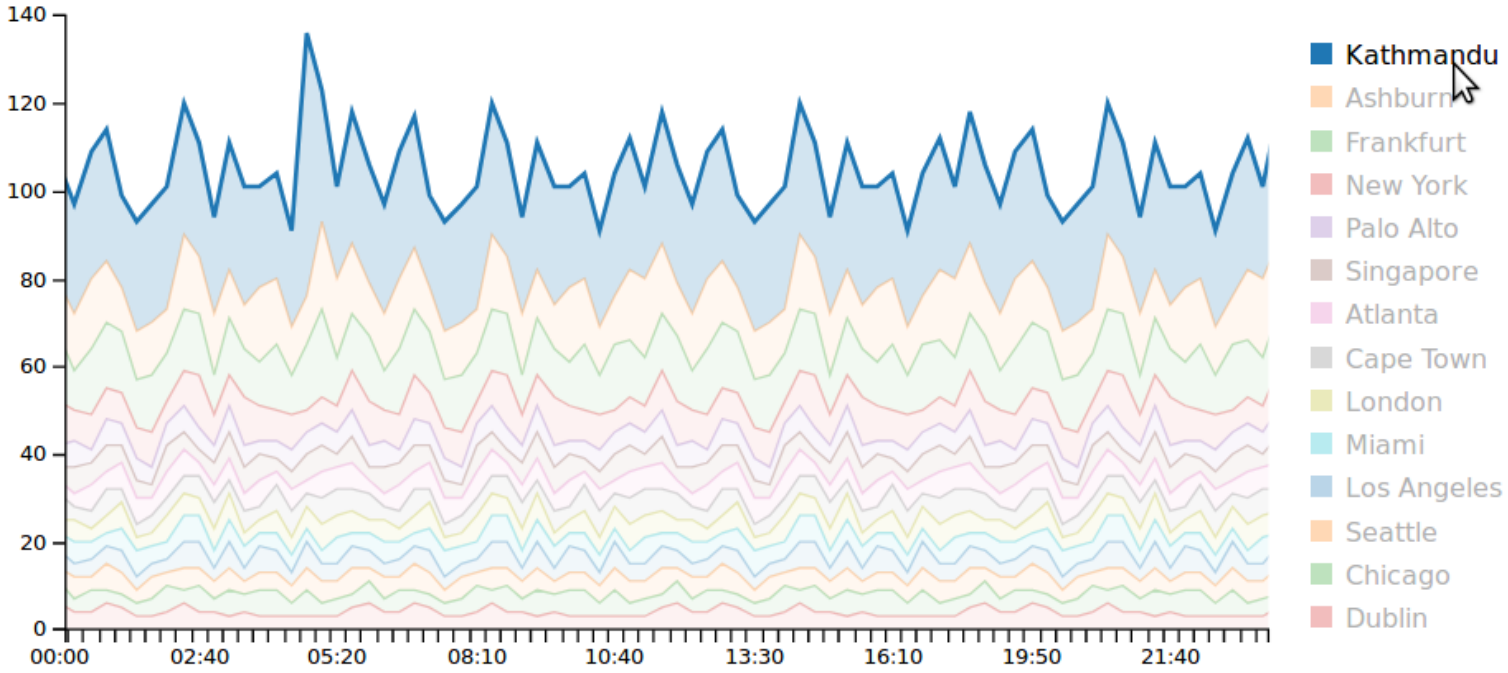

PCH
Packet Clearing House

Logged in as **mrjones**

PCH
Packet Clearing House

Admin | About | Services | Tools | Resources

Home > Account > DNS Stats Graph

DNS Stats - .aj - Last 24 hours

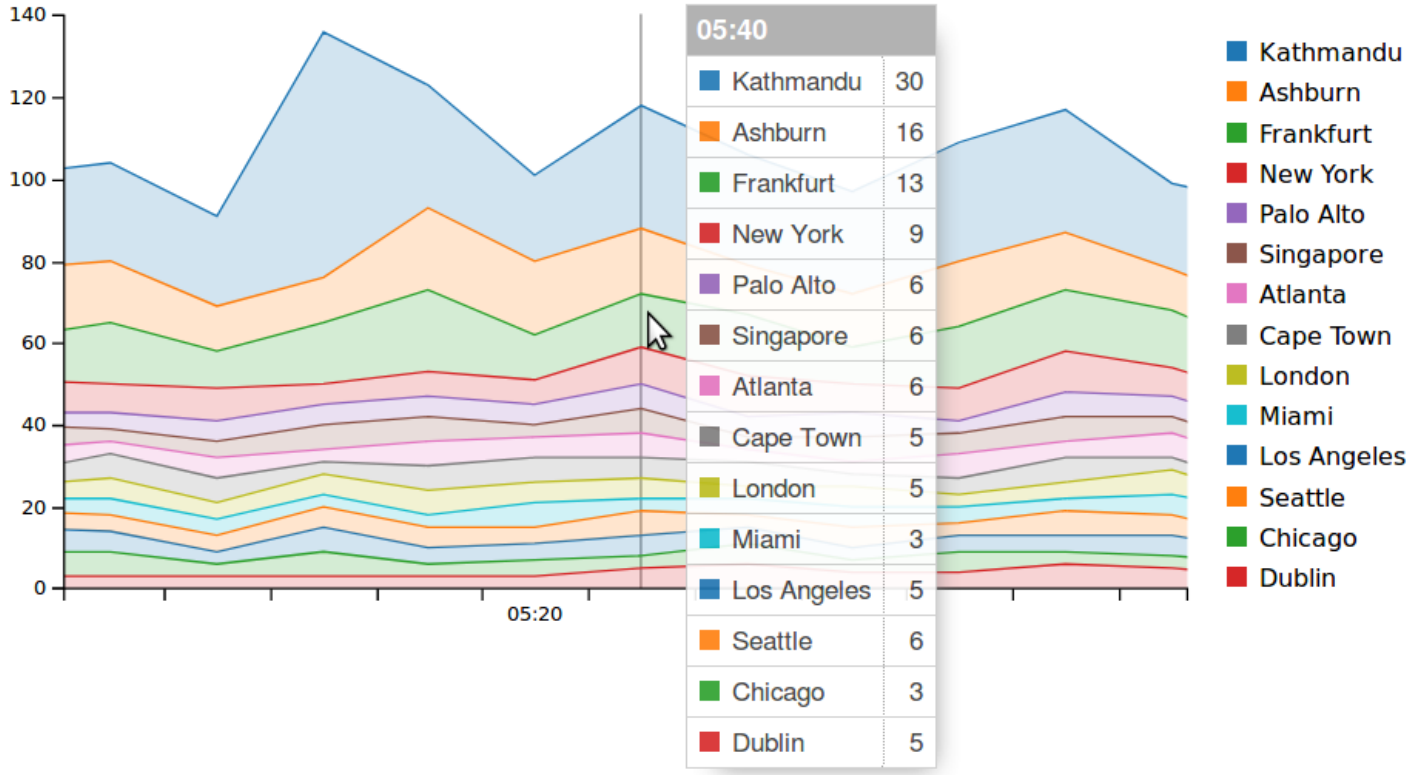

DNS Stats - .aj - Current vs Baseline By City

Download

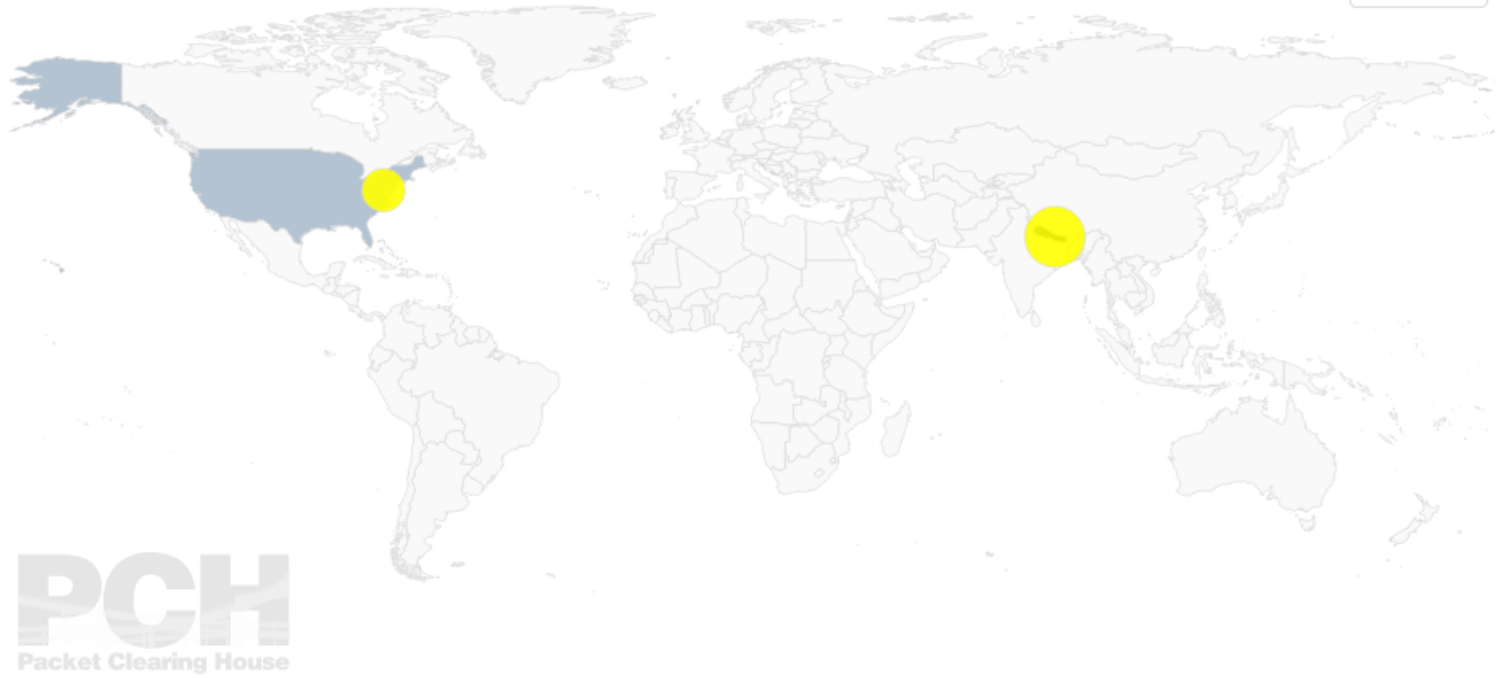

Each marker represents a city. Color of country represents the sum of queries in that country. Source: **pch.net**

Filters

Queries is > 900000 – Remove this filter

+ New filter

		QPS		TCP %		IPv6 %	
City ▲	Country	Current	Baseline	Current	Baseline	Current	Baseline
Ashburn	US	9246 (5%)	2788 (31%)	2%	4%	5%	5%
Kathmandu	NP	8729 (1%)	2661 (41%)	2%	4%	1%	4%

PCH
Packet Clearing House

DNS Stats - .aj - Current vs Baseline By City

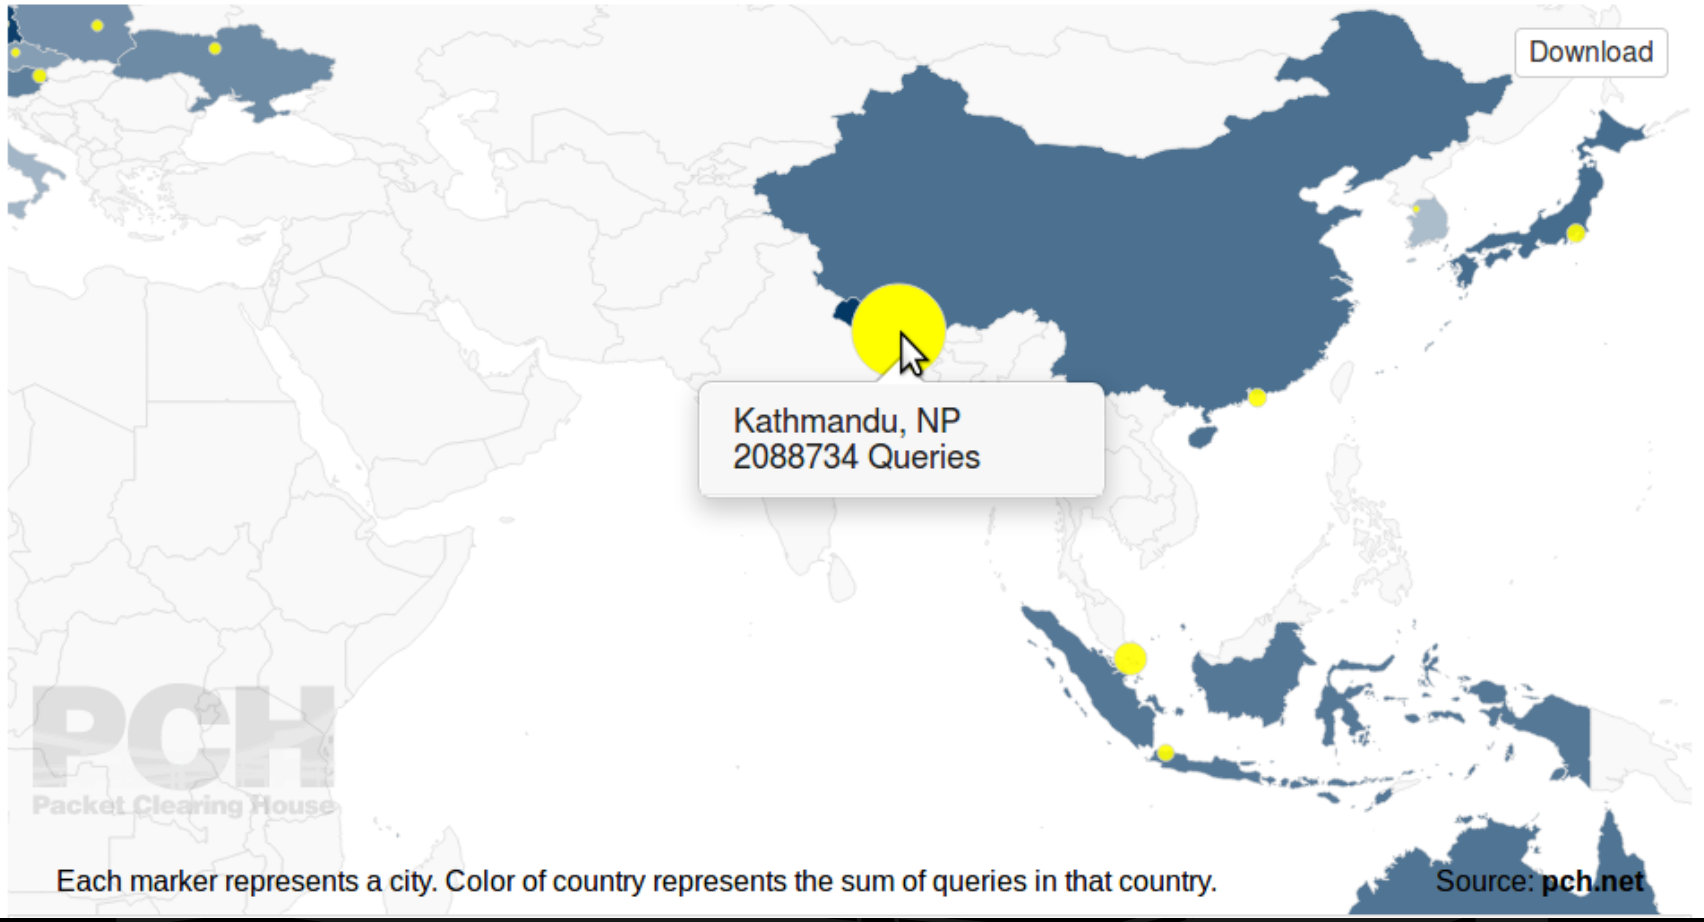

PCH
Packet Clearing House

DNS Stats - .aj - Current vs Baseline By City

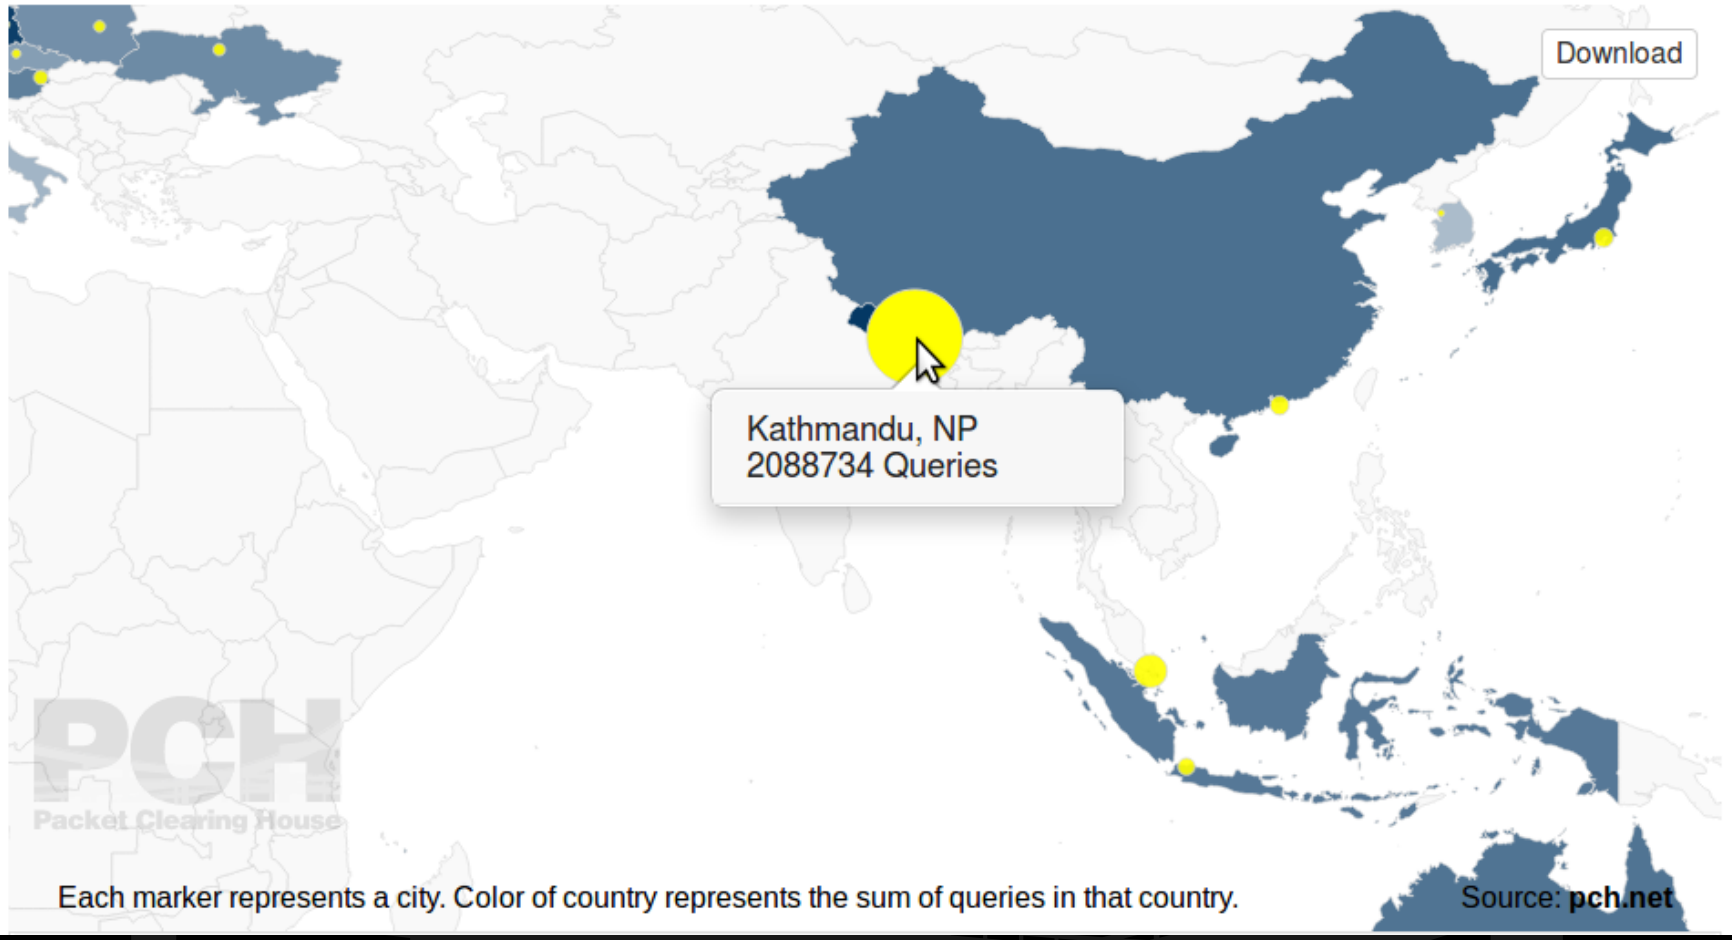

PCH
Packet Clearing House

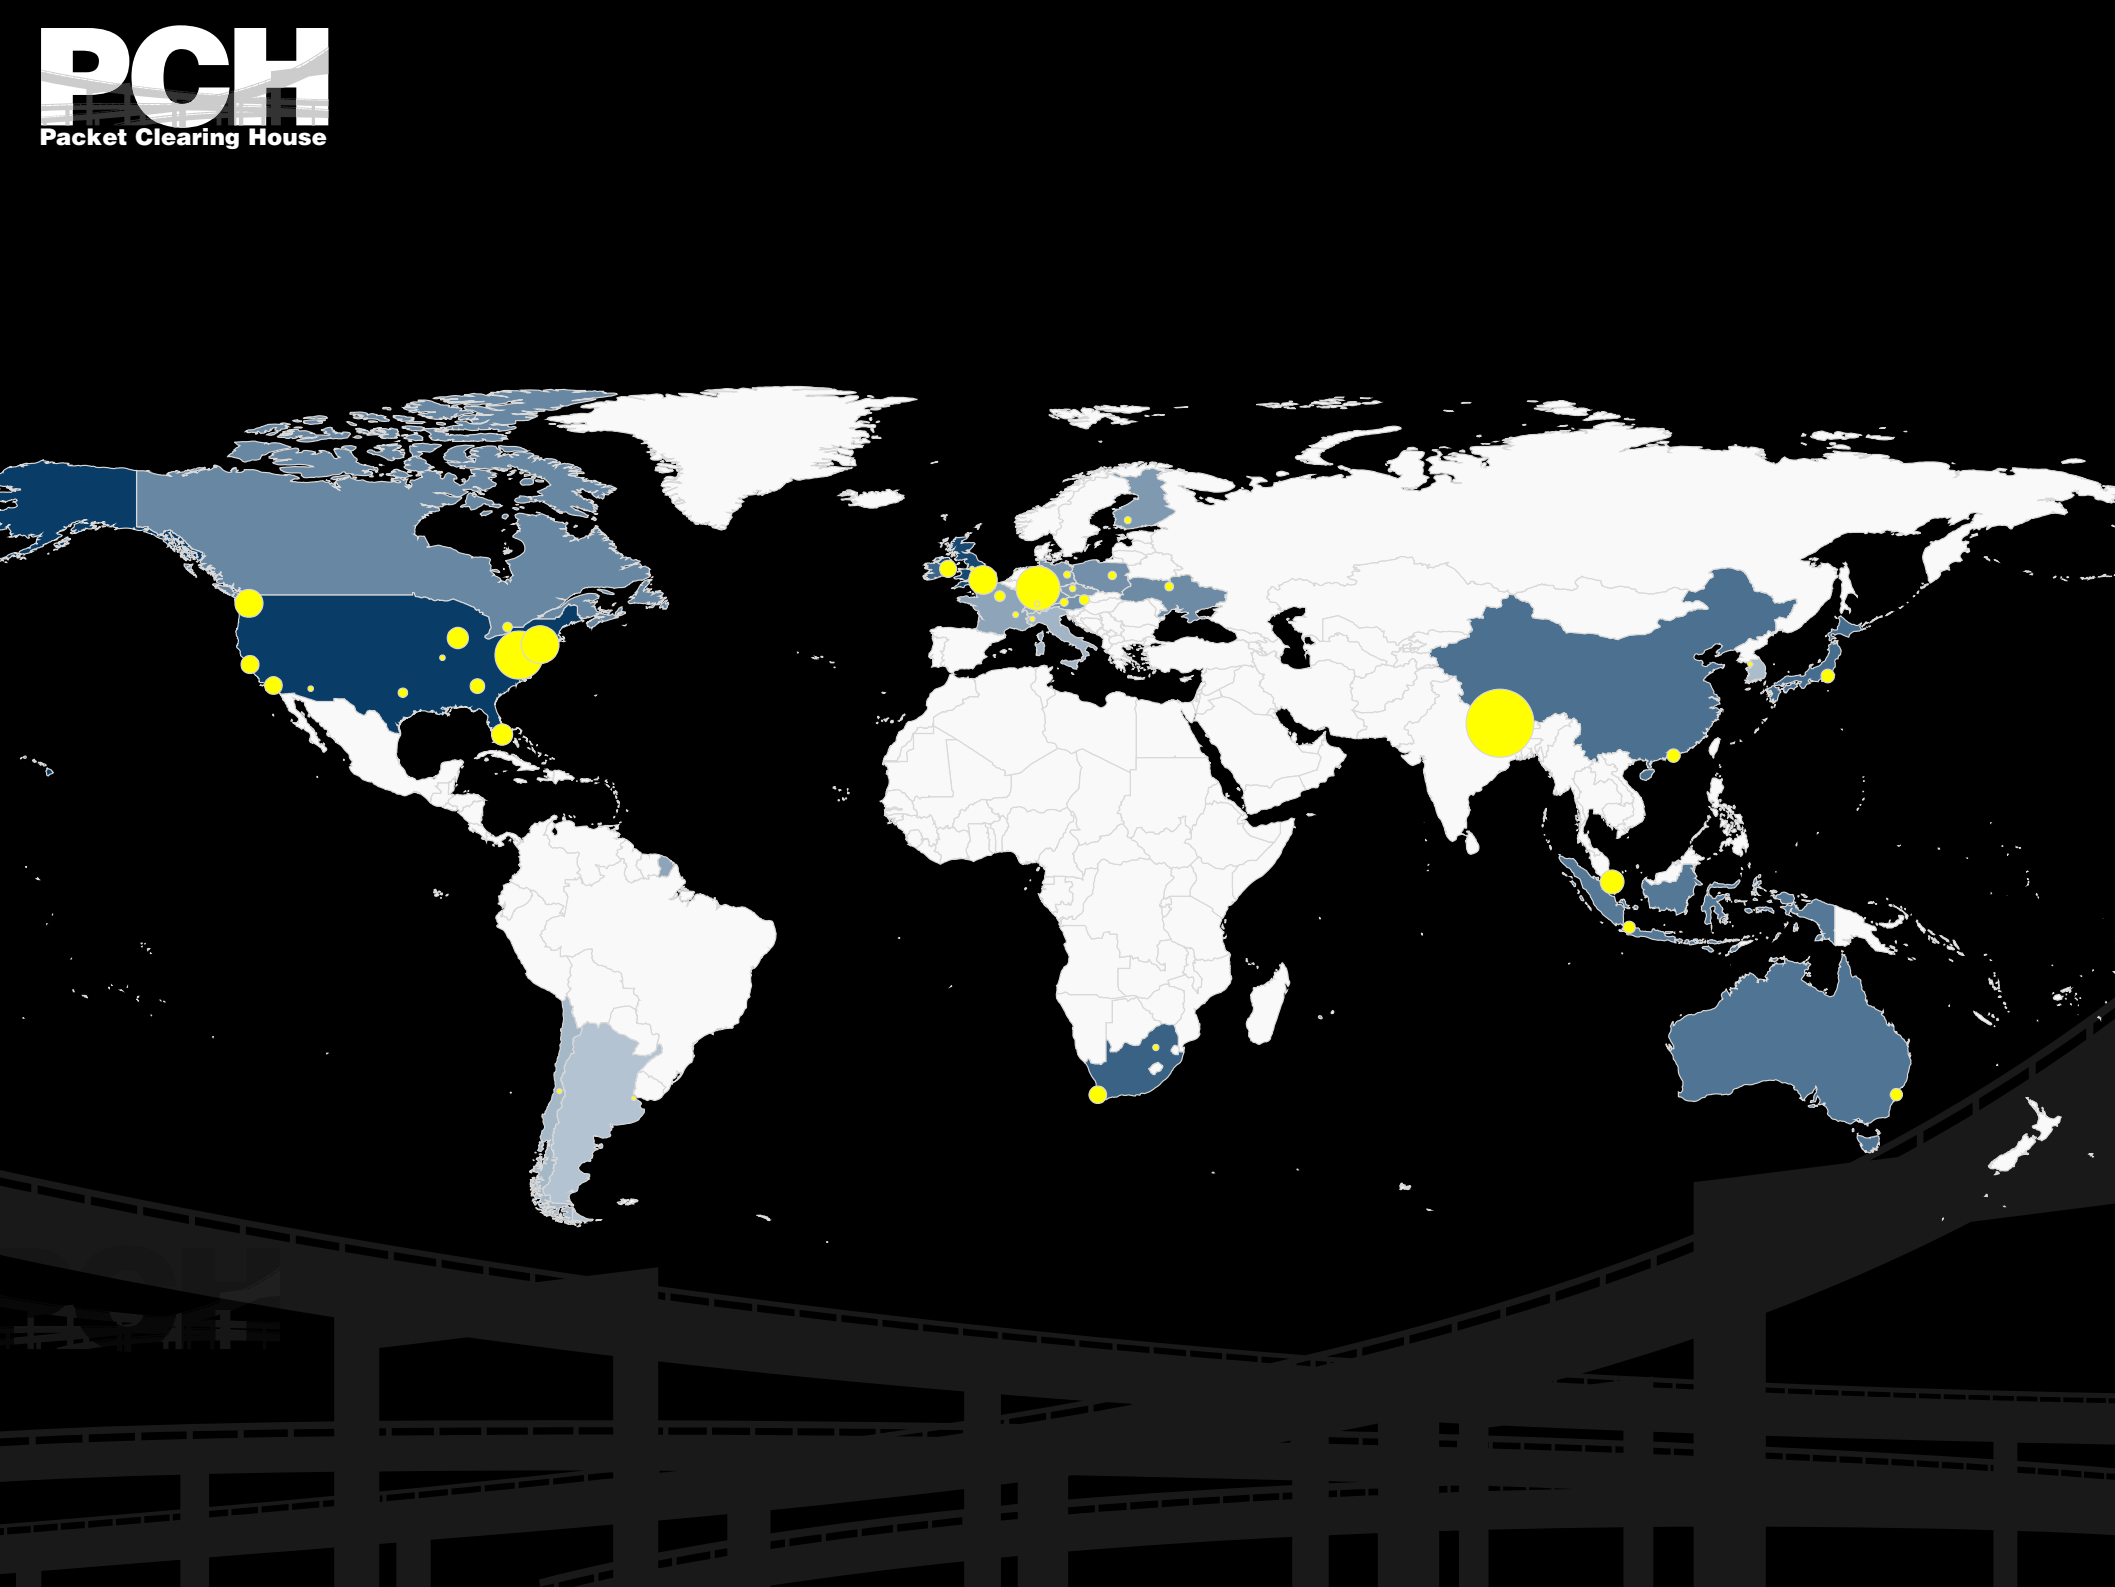

PCH
Packet Clearing House

Psuedo & Sample Code

PCH
Packet Clearing House

MapTable

```
<div id='vizContainer'></div>
<script src="d3.min.js"></script>
<script src="topojson.min.js"></script>
<script src="maptable.min.js"></script>
<script>
var viz = d3.maptable('#vizContainer')
  .csv('/api/dns_stats.csv?tld=aj')
  .map({path:'/maps/world110m.json'})
  .filters()
  .table()
  .render();
</script>
```

PCH
Packet Clearing House

InfluxDB 

```
$client = new IfDB\Client($host,
$port);
$db = IfDB\Client::fromDSN(sprintf(
   'influxdb://user:pass@%s:%s/%s',
   $host, $port, $dbname));
$client = $database->getClient();
```

PCH
Packet Clearing House

InfluxDB 

```
$result = $db->getQueryBuilder()
   ->select('tcp_count, udp_count,
etc')
   ->from('dns_stats')
   ->where(["tld = 'aj'"])
   ->getPoints();
```

PCH
Packet Clearing House

DNSAuth

```
cp mon-01.sample.net_2017-10-17.17-07.dmp.gz
/dnsauth/ingest

sudo ./go/bin/DNSAuth -c dnsauth/DNSAuth/dnsauth.toml

2017/12/12 06:55:46 Loading config file…
2017/12/12 06:55:46 OK!
2017/12/12 06:55:46 Getting customer list from
postgres…
2017/12/12 06:55:46 OK!
2017/12/12 06:56:16 Processed dump [mon-01](2017-10-
17 17:07:00 +0000 UTC - 2017-10-17 17:10:00.215724
+0000 UTC): 833 lines in (2.876312ms) seconds!
```

PCH
Packet Clearing House

c3

```
var chart = c3.generate({
    bindto: '#chart',
    data: {
        order: 'desc',
        columns: [ ['SET1',...],],
        types: {'SET1': 'area',... },
        groups: [['SET1',...]],
    },
    axis: {
        X: {
            type: 'category',
            Categories: [ X_LABELS_HERE],
        },},});
```

PCH
Packet Clearing House

c3

```
data: {
    order: 'desc',
    columns: [
        ['Kathmandu',27,25,29,30,21,25,27]
        ['Ashburn',12,13,16,14,10,11,12,10]
    ,],
    Types: {
        'Kathmandu': 'area', 'Ashburn': 'area',

    },
    Groups: [
        'Kathmandu', 'Ashburn',
    ],
},
```

PCH
Packet Clearing House

c3

```
axis: {
   X: {
      type: 'category',
      Categories: [ '00:00', '00:20',
'00:40', '00:50', '01:10',  '01:30',
'01:40', '02:10', '02:20', '02:40',
'03:00',],
    },
}
```