pktvisor

open source packet stream summarizer

Shannon Weyrick, VP Architecture @ NS1
sweyrick@ns1.com

v3

What is pktvisor?

- Open source **network visibility** tool
- **Summarizes** traffic in real time at edges with data sketches
- Includes a **command line interface** for on-node visualization
- Includes an **HTTP API** for collecting summaries to a central location
- Metrics include
  - Packet counts and rates (w/percentiles), breakdown by ingress/egress, protocol
  - DNS counts and rates, breakdown by protocol, response code
  - Cardinality: Source and destination IP, DNS Qname
  - DNS transaction timings (w/percentiles)
  - Top 10 heavy hitters for
    - IPs and ports
    - DNS Qnames, Qtypes, Result Codes
    - Slow DNS transactions, NX, SRVFAIL, REFUSED Qnames
    - GeoIP and ASN

NS1.

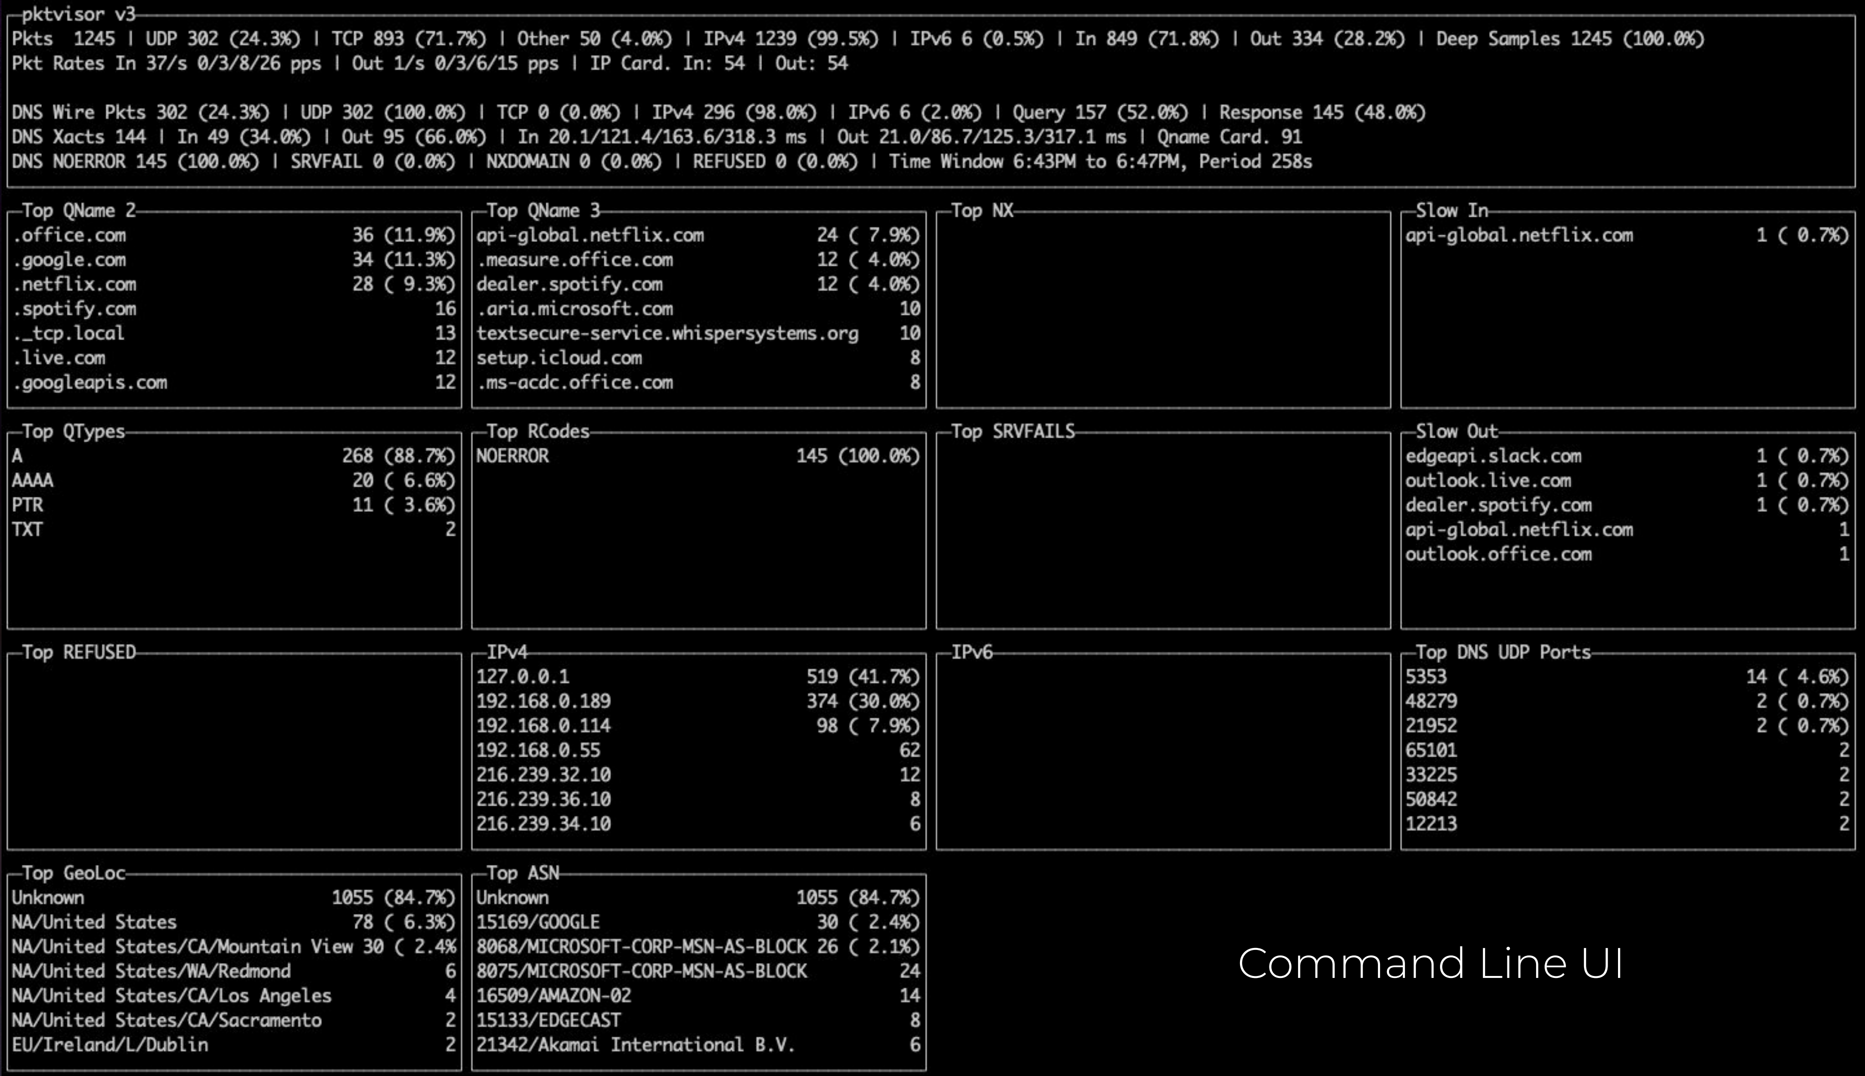

```
┌pktvisor v3──────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────┐
│Pkts  1245 | UDP 302 (24.3%) | TCP 893 (71.7%) | Other 50 (4.0%) | IPv4 1239 (99.5%) | IPv6 6 (0.5%) | In 849 (71.8%) | Out 334 (28.2%) | Deep Samples 1245 (100.0%)│
│Pkt Rates In 37/s 0/3/8/26 pps | Out 1/s 0/3/6/15 pps | IP Card. In: 54 | Out: 54                                                                      │
│                                                                                                                                                        │
│DNS Wire Pkts 302 (24.3%) | UDP 302 (100.0%) | TCP 0 (0.0%) | IPv4 296 (98.0%) | IPv6 6 (2.0%) | Query 157 (52.0%) | Response 145 (48.0%)               │
│DNS Xacts 144 | In 49 (34.0%) | Out 95 (66.0%) | In 20.1/121.4/163.6/318.3 ms | Out 21.0/86.7/125.3/317.1 ms | Qname Card. 91                         │
│DNS NOERROR 145 (100.0%) | SRVFAIL 0 (0.0%) | NXDOMAIN 0 (0.0%) | REFUSED 0 (0.0%) | Time Window 6:43PM to 6:47PM, Period 258s                          │
└────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────────┘
┌Top QName 2──────────────────────┐┌Top QName 3────────────────────────────┐┌Top NX───────────────────────────┐┌Slow In──────────────────────────┐
│.office.com          36 (11.9%)│││api-global.netflix.com        24 ( 7.9%)││                                 ││api-global.netflix.com  1 ( 0.7%)│
│.google.com          34 (11.3%)│││.measure.office.com           12 ( 4.0%)││                                 ││                                 │
│.netflix.com         28 ( 9.3%)│││dealer.spotify.com            12 ( 4.0%)││                                 ││                                 │
│.spotify.com                 16│││.aria.microsoft.com                   10││                                 ││                                 │
│._tcp.local                  13│││textsecure-service.whispersystems.org 10││                                 ││                                 │
│.live.com                    12│││setup.icloud.com                       8││                                 ││                                 │
│.googleapis.com              12│││.ms-acdc.office.com                    8││                                 ││                                 │
└─────────────────────────────────┘└───────────────────────────────────────┘└─────────────────────────────────┘└─────────────────────────────────┘
┌Top QTypes───────────────────────┐┌Top RCodes─────────────────────────────┐┌Top SRVFAILS─────────────────────┐┌Slow Out─────────────────────────┐
│A                   268 (88.7%)  ││NOERROR                    145 (100.0%)││                                 ││edgeapi.slack.com       1 ( 0.7%)│
│AAAA                 20 ( 6.6%)  ││                                       ││                                 ││outlook.live.com        1 ( 0.7%)│
│PTR                  11 ( 3.6%)  ││                                       ││                                 ││dealer.spotify.com      1 ( 0.7%)│
│TXT                           2  ││                                       ││                                 ││api-global.netflix.com          1│
│                                 ││                                       ││                                 ││outlook.office.com              1│
└─────────────────────────────────┘└───────────────────────────────────────┘└─────────────────────────────────┘└─────────────────────────────────┘
┌Top REFUSED──────────────────────┐┌IPv4───────────────────────────────────┐┌IPv6─────────────────────────────┐┌Top DNS UDP Ports────────────────┐
│                                 ││127.0.0.1                   519 (41.7%)││                                 ││5353                   14 ( 4.6%)│
│                                 ││192.168.0.189               374 (30.0%)││                                 ││48279                   2 ( 0.7%)│
│                                 ││192.168.0.114                98 ( 7.9%)││                                 ││21952                   2 ( 0.7%)│
│                                 ││192.168.0.55                         62││                                 ││65101                           2│
│                                 ││216.239.32.10                        12││                                 ││33225                           2│
│                                 ││216.239.36.10                         8││                                 ││50842                           2│
│                                 ││216.239.34.10                         6││                                 ││12213                           2│
└─────────────────────────────────┘└───────────────────────────────────────┘└─────────────────────────────────┘└─────────────────────────────────┘
┌Top GeoLoc───────────────────────────────────┐┌Top ASN──────────────────────────────────────┐
│Unknown                          1055 (84.7%)││Unknown                          1055 (84.7%)│
│NA/United States                   78 ( 6.3%)││15169/GOOGLE                       30 ( 2.4%)│
│NA/United States/CA/Mountain View 30 ( 2.4%  ││8068/MICROSOFT-CORP-MSN-AS-BLOCK 26 ( 2.1%)  │
│NA/United States/WA/Redmond                 6││8075/MICROSOFT-CORP-MSN-AS-BLOCK           24│
│NA/United States/CA/Los Angeles             4││16509/AMAZON-02                            14│
│NA/United States/CA/Sacramento              2││15133/EDGECAST                              8│
│EU/Ireland/L/Dublin                         2││21342/Akamai International B.V.             6│
└─────────────────────────────────────────────┘└─────────────────────────────────────────────┘
```

Command Line UI

Motivation; pktvisor v1

- 2014, needed more **visibility** across our **global anycast network**
- Nominal operations, debugging, DDoS
- Forked netsniff-ng to make pktvisor v1 (remains open source)
- Essentially a **DNS "top"** on node with CLI UI
- We did some automated central collection, mixed results
- Problems
    - resource usage (did not use sketches)
    - missing IPv6 and TCP support
    - hard to collect centrally
    - did not track transactions (query/reply pair)
    - each process ran a new analyzer

pktvisor Rewrite Guiding Principles

1. **Summarize, don’t collect**
    a. we are interested in a distilled signal, not the raw stream
    b. localized but real-time at the source
    c. global view but some lag centrally
2. **Sliding time window, JSON interface**
    a. maintain only easily consumed summary of N minutes of data
3. **Plug the v1 holes, Enhance**
    a. IPv6, TCP, efficiency, new metrics

Goal: Summarize, don’t collect

- This solution is **not for data warehousing**
  - it will not provide an “audit log” of all packet information
- Instead, **summarize** with counters and sketches directly at edge
- Reduces complexity, at the expense of querying flexibility
  - lightweight data requirements, less complex distributed system
  - but you cannot ask it arbitrary questions on raw data
- Makes for fast dashboards (local and central), with **low network and storage requirements**
- ~7kb JSON per 1 minute summary, per host
  - 100 hosts generate <1Mb per min == ~1 Gb per day (uncompressed)
  - data rate is a function of the number of hosts, *not a function of traffic rates*
  - traffic spikes and DDoS do not affect downstream collection systems

Goal: Sliding time window, JSON interface

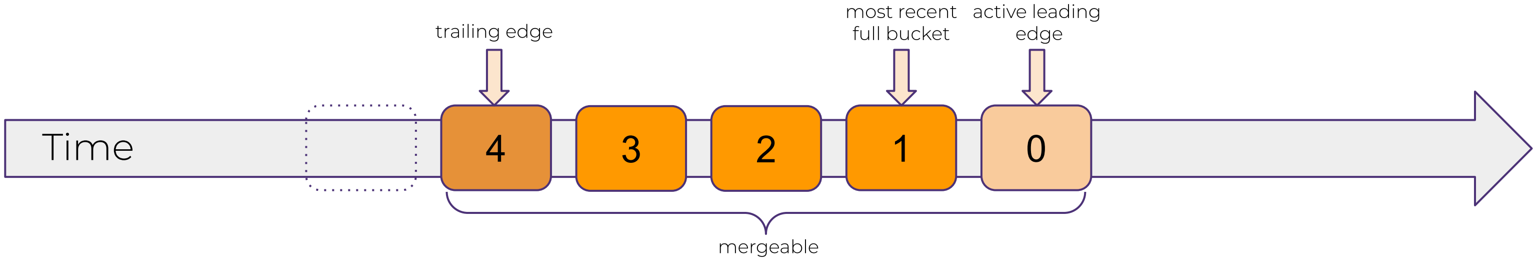

- Maintains N individual mins (default 5) of metrics which may be **merged** to provide summary across full window
- Always-on daemon supplies information to CLI UI *and* central collection via HTTP
- Both merged and individual minute buckets are available for collection in **REST API**
    - CLI UI uses the merged window
    - Central collector gathers a single minute, once a minute
- Not opinionated on which collector/central database is used.

NS1.

Goal: Plug the v1 holes, Enhance

- IPv6 and TCP fully supported
- Local CLI UI can efficiently run multiple times on the same node, or connect remotely
- Tracks **DNS transactions** (query/reply pairs), in and out
- Operate on a **pcap** in addition to **live capture**
  - pcap will summarize JSON to stdout
- Sketches allow new metrics: **cardinality, quantiles**
- Adjustable deep **sample rate**
  - deep sample invokes full L7 parse and full data sketch update
  - without deep sample, only simple counters are updated

NS1.

Under the hood

- Main capture daemon is **written in C++**
- **CLI UI** is in written in **golang** (UI is gocui)
- PcapPlusPlus, libpcap based but PF_RING & DPDK possibile
- Apache DataSketches
- Optional MaxMind support for GeoIP and ASN
- Small code base and mostly header only libraries
- Docker first (available on Dockerhub), easy to try out
- Performance numbers for 3.0.7 (single instance)
  - low load ~13 MB resident ram
  - fully production loaded 5m window == ~200 MB resident ram
  - >100k QPS live capture before packet buffer drops seen
  - expecting several paths for optimization

DataSketches

- Relies on Apache DataSketches
- **Fast, probabilistic** data structures designed for streaming
- Results are **approximate** but within **well defined error bounds**
- Provides **cardinality, heavy hitters (frequent items), quantiles**
- Designed to be **merged**, which is how we support time window
- Possible to expose the raw binary sketch data in the API so that it can be **merged across hosts and data centers**

NS1.

Command Line UI

NS1.

Command Line UI

- **Visualize all stats in entire time window** on a **single node** in real time
- May connect to local or remote node
- Multiple operators may efficiently visualize at the same time
- Updates results 1 per second
    - reminiscent of command line “top”

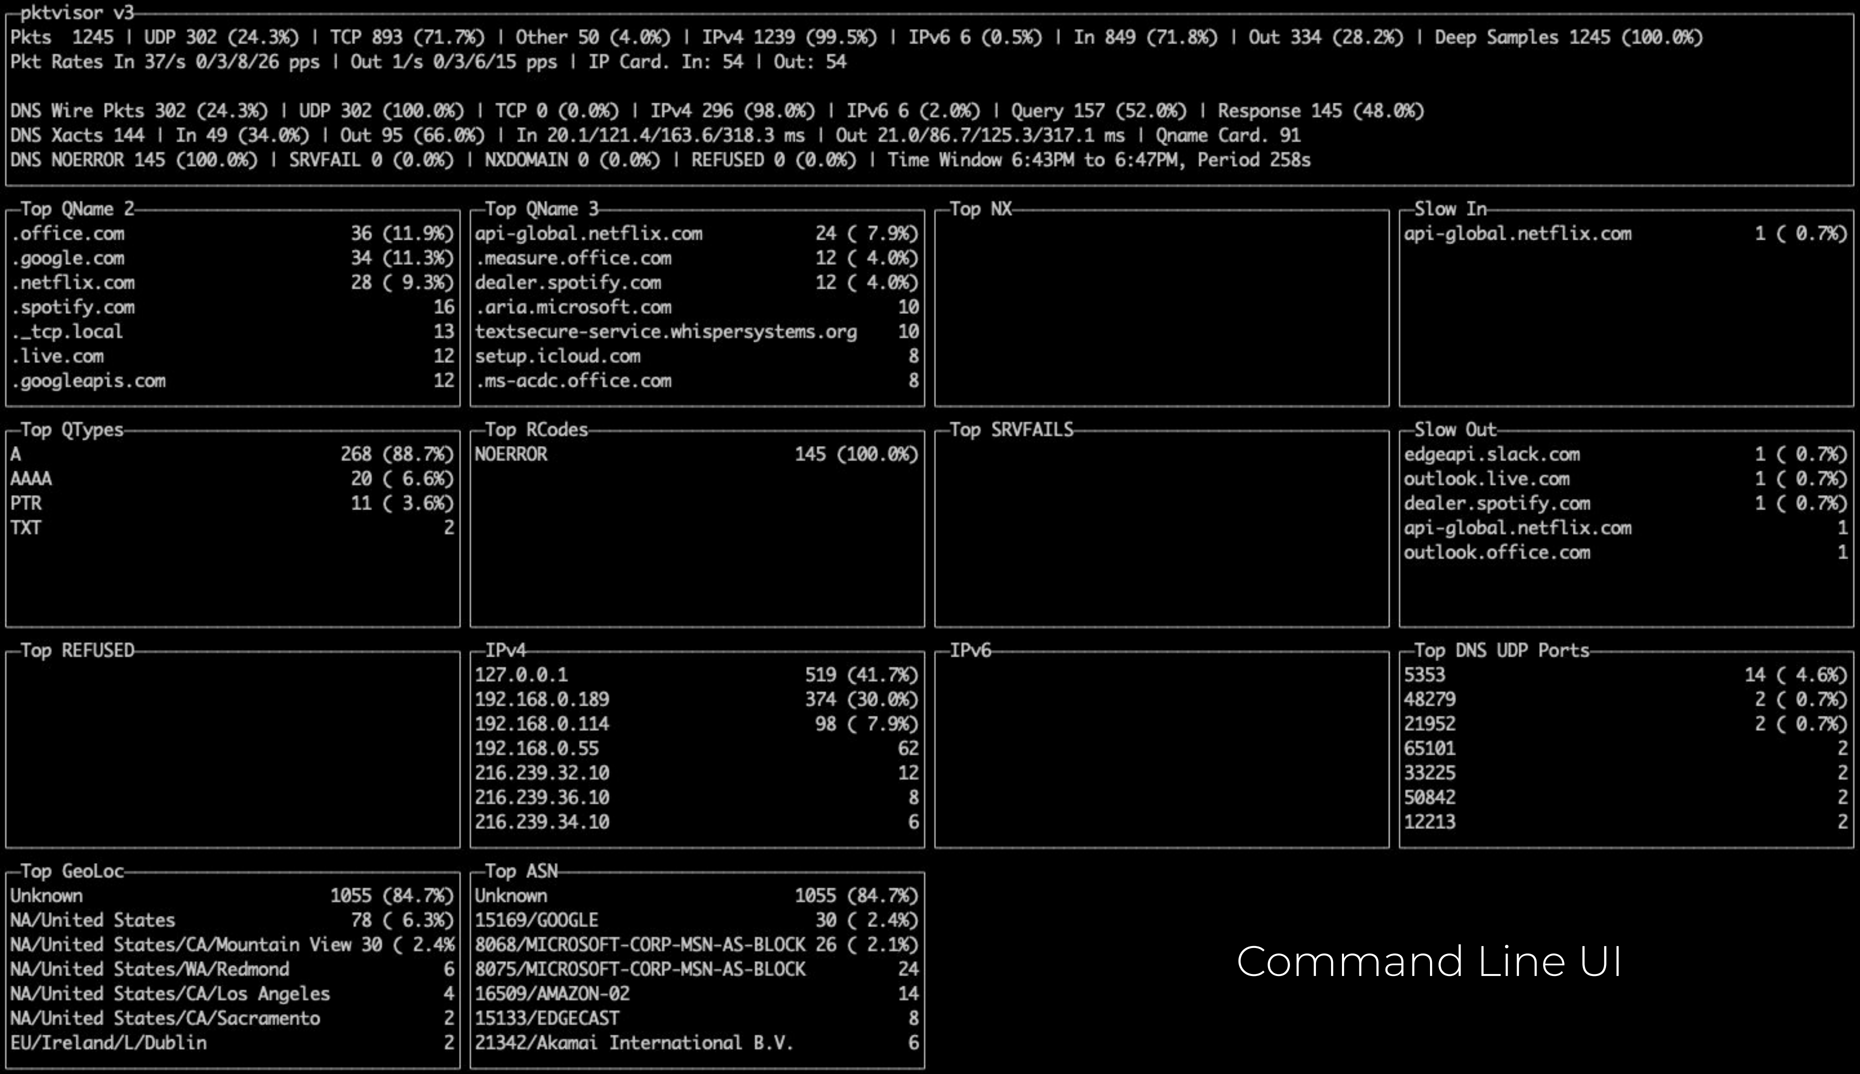

```
pktvisor v3
Pkts  1245 | UDP 302 (24.3%) | TCP 893 (71.7%) | Other 50 (4.0%) | IPv4 1239 (99.5%) | IPv6 6 (0.5%) | In 849 (71.8%) | Out 334 (28.2%) | Deep Samples 1245 (100.0%)
Pkt Rates In 37/s 0/3/8/26 pps | Out 1/s 0/3/6/15 pps | IP Card. In: 54 | Out: 54

DNS Wire Pkts 302 (24.3%) | UDP 302 (100.0%) | TCP 0 (0.0%) | IPv4 296 (98.0%) | IPv6 6 (2.0%) | Query 157 (52.0%) | Response 145 (48.0%)
DNS Xacts 144 | In 49 (34.0%) | Out 95 (66.0%) | In 20.1/121.4/163.6/318.3 ms | Out 21.0/86.7/125.3/317.1 ms | Qname Card. 91
DNS NOERROR 145 (100.0%) | SRVFAIL 0 (0.0%) | NXDOMAIN 0 (0.0%) | REFUSED 0 (0.0%) | Time Window 6:43PM to 6:47PM, Period 258s

Top QName 2                          Top QName 3                                    Top NX        Slow In
.office.com            36 (11.9%)    api-global.netflix.com            24 ( 7.9%)                 api-global.netflix.com     1 ( 0.7%)
.google.com            34 (11.3%)    .measure.office.com               12 ( 4.0%)
.netflix.com           28 ( 9.3%)    dealer.spotify.com                12 ( 4.0%)
.spotify.com                   16    .aria.microsoft.com                       10
._tcp.local                    13    textsecure-service.whispersystems.org     10
.live.com                      12    setup.icloud.com                           8
.googleapis.com                12    .ms-acdc.office.com                        8

Top QTypes                           Top RCodes                                     Top SRVFAILS  Slow Out
A                     268 (88.7%)    NOERROR                          145 (100.0%)                edgeapi.slack.com          1 ( 0.7%)
AAAA                   20 ( 6.6%)                                                                 outlook.live.com           1 ( 0.7%)
PTR                    11 ( 3.6%)                                                                 dealer.spotify.com         1 ( 0.7%)
TXT                             2                                                                 api-global.netflix.com            1
                                                                                                  outlook.office.com                1

Top REFUSED                          IPv4                                           IPv6          Top DNS UDP Ports
                                     127.0.0.1                        519 (41.7%)                 5353                      14 ( 4.6%)
                                     192.168.0.189                    374 (30.0%)                 48279                      2 ( 0.7%)
                                     192.168.0.114                     98 ( 7.9%)                 21952                      2 ( 0.7%)
                                     192.168.0.55                              62                 65101                             2
                                     216.239.32.10                             12                 33225                             2
                                     216.239.36.10                              8                 50842                             2
                                     216.239.34.10                              6                 12213                             2

Top GeoLoc                                          Top ASN
Unknown                               1055 (84.7%)  Unknown                                1055 (84.7%)
NA/United States                        78 ( 6.3%)  15169/GOOGLE                             30 ( 2.4%)
NA/United States/CA/Mountain View 30 ( 2.4%         8068/MICROSOFT-CORP-MSN-AS-BLOCK 26 ( 2.1%)
NA/United States/WA/Redmond                      6  8075/MICROSOFT-CORP-MSN-AS-BLOCK                 24
NA/United States/CA/Los Angeles                  4  16509/AMAZON-02                                  14
NA/United States/CA/Sacramento                   2  15133/EDGECAST                                    8
EU/Ireland/L/Dublin                              2  21342/Akamai International B.V.                   6
```

Command Line UI

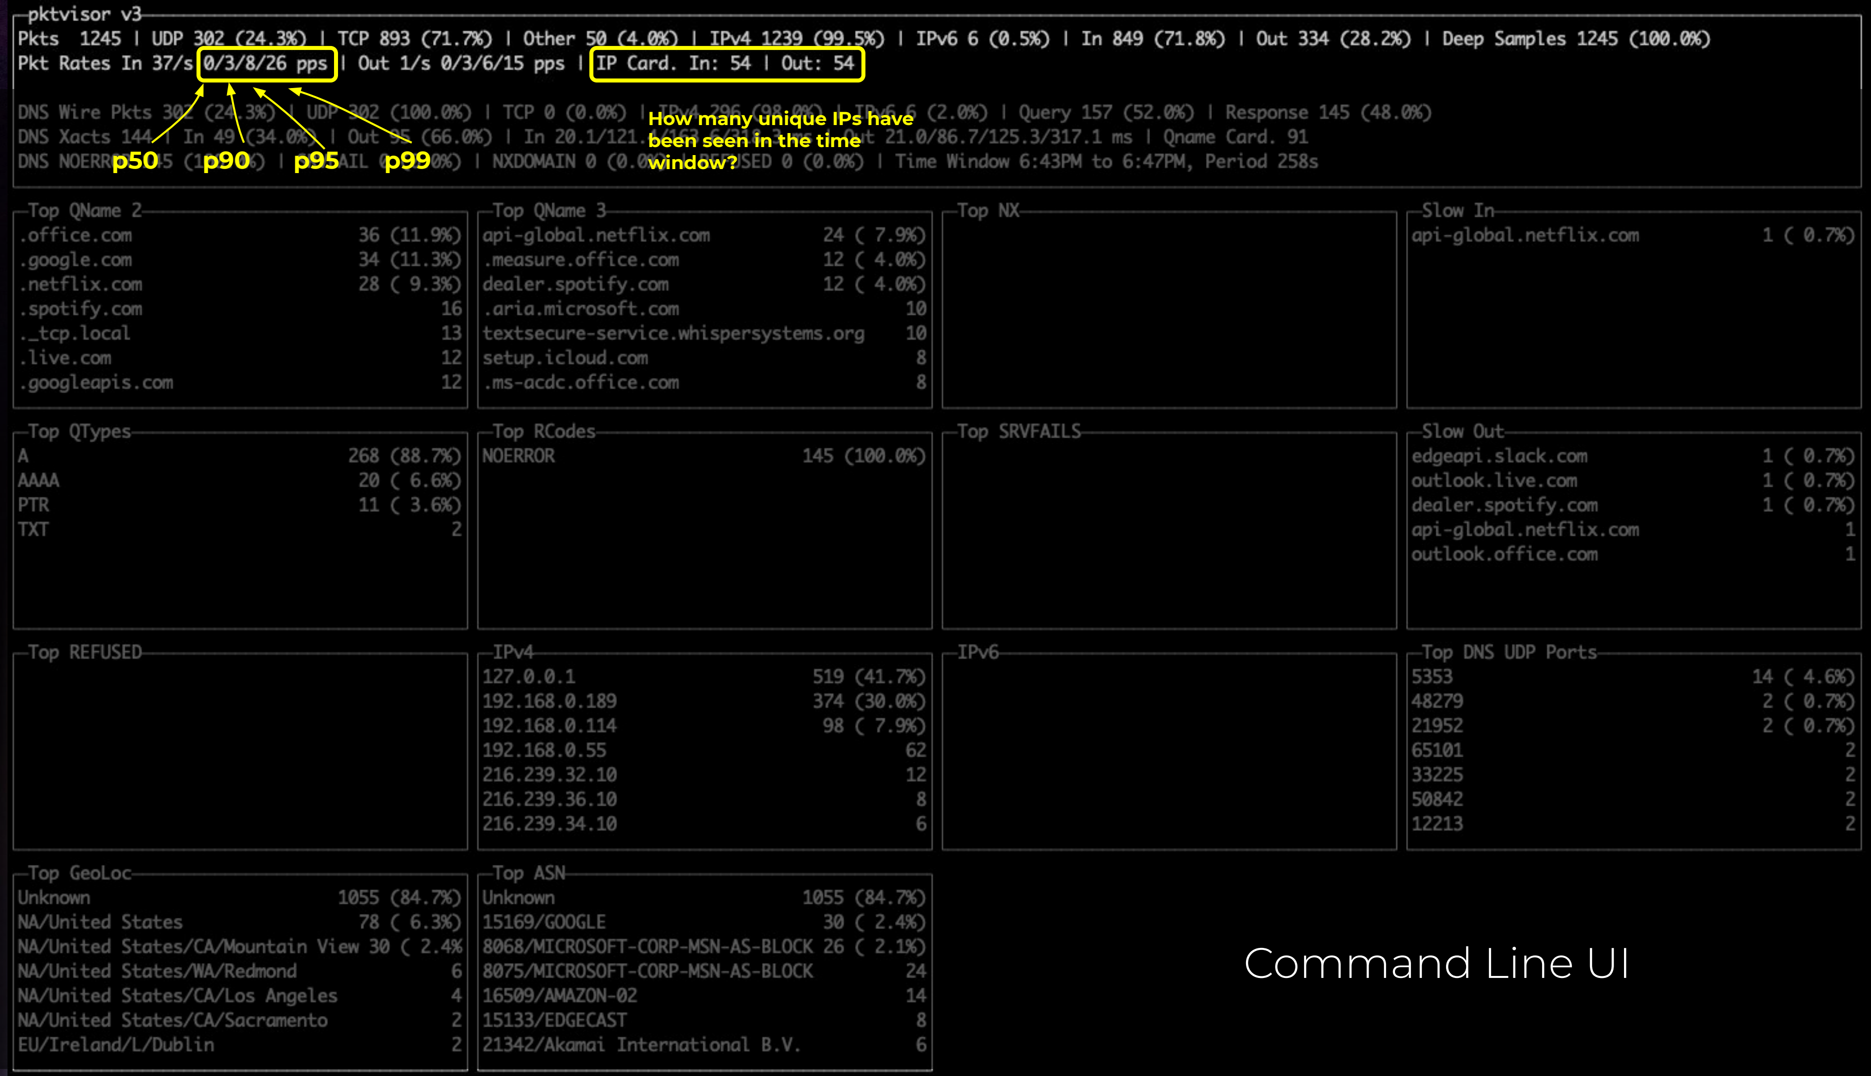

```
pktvisor v3
Pkts  1245 | UDP 302 (24.3%) | TCP 893 (71.7%) | Other 50 (4.0%) | IPv4 1239 (99.5%) | IPv6 6 (0.5%) | In 849 (71.8%) | Out 334 (28.2%) | Deep Samples 1245 (100.0%)
Pkt Rates In 37/s 0/3/8/26 pps | Out 1/s 0/3/6/15 pps | IP Card. In: 54 | Out: 54

DNS Wire Pkts 302 (24.3%) | UDP 302 (100.0%) | TCP 0 (0.0%) | IPv4 296 (98.0%) | IPv6 6 (2.0%) | Query 157 (52.0%) | Response 145 (48.0%)
DNS Xacts 144 | In 49 (34.0%) | Out 95 (66.0%) | In 20.1/121.4/163.6/318.3 ms | Out 21.0/86.7/125.3/317.1 ms | Qname Card. 91
DNS NOERROR 145 (100.0%) | SRVFAIL 0 (0.0%) | NXDOMAIN 0 (0.0%) | REFUSED 0 (0.0%) | Time Window 6:43PM to 6:47PM, Period 258s

Top QName 2                          Top QName 3                                   Top NX        Slow In
.office.com           36 (11.9%)     api-global.netflix.com           24 ( 7.9%)                 api-global.netflix.com        1 ( 0.7%)
.google.com           34 (11.3%)     .measure.office.com              12 ( 4.0%)
.netflix.com          28 ( 9.3%)     dealer.spotify.com               12 ( 4.0%)
.spotify.com                  16     .aria.microsoft.com                      10
._tcp.local                   13     textsecure-service.whispersystems.org    10
.live.com                     12     setup.icloud.com                          8
.googleapis.com               12     .ms-acdc.office.com                       8

Top QTypes                           Top RCodes                                    Top SRVFAILS  Slow Out
A                    268 (88.7%)     NOERROR                         145 (100.0%)                edgeapi.slack.com             1 ( 0.7%)
AAAA                  20 ( 6.6%)                                                                 outlook.live.com              1 ( 0.7%)
PTR                   11 ( 3.6%)                                                                 dealer.spotify.com            1 ( 0.7%)
TXT                            2                                                                 api-global.netflix.com                1
                                                                                                 outlook.office.com                    1

Top REFUSED                          IPv4                                          IPv6          Top DNS UDP Ports
                                     127.0.0.1                       519 (41.7%)                 5353                         14 ( 4.6%)
                                     192.168.0.189                   374 (30.0%)                 48279                         2 ( 0.7%)
                                     192.168.0.114                    98 ( 7.9%)                 21952                         2 ( 0.7%)
                                     192.168.0.55                             62                 65101                                 2
                                     216.239.32.10                            12                 33225                                 2
                                     216.239.36.10                             8                 50842                                 2
                                     216.239.34.10                             6                 12213                                 2

Top GeoLoc                                   Top ASN
Unknown                         1055 (84.7%) Unknown                                  1055 (84.7%)
NA/United States                  78 ( 6.3%) 15169/GOOGLE                               30 ( 2.4%)
NA/United States/CA/Mountain View 30 ( 2.4%  8068/MICROSOFT-CORP-MSN-AS-BLOCK           26 ( 2.1%)
NA/United States/WA/Redmond                6 8075/MICROSOFT-CORP-MSN-AS-BLOCK                   24
NA/United States/CA/Los Angeles            4 16509/AMAZON-02                                    14
NA/United States/CA/Sacramento             2 15133/EDGECAST                                      8
EU/Ireland/L/Dublin                        2 21342/Akamai International B.V.                     6
```

p50

p90

p95

p99

How many unique IPs have been seen in the time window?

Command Line UI

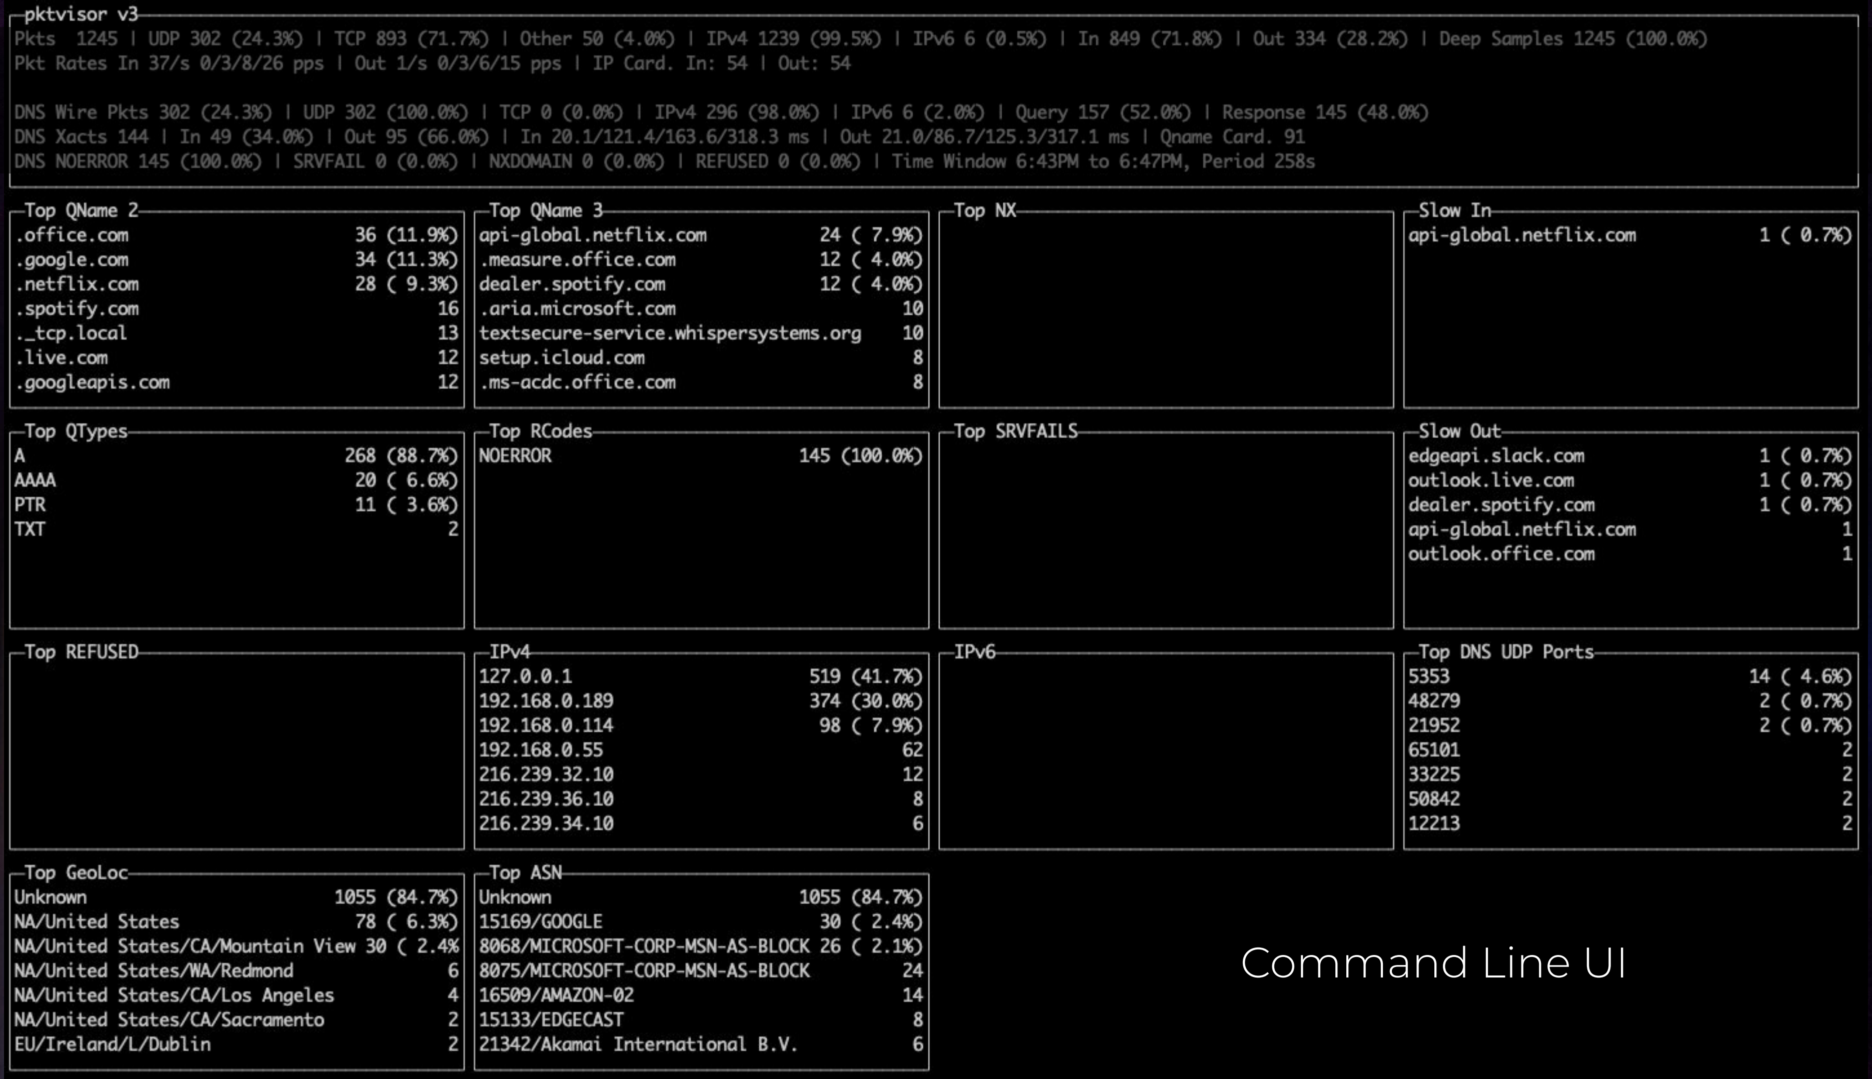

```
pktvisor v3
Pkts  1245 | UDP 302 (24.3%) | TCP 893 (71.7%) | Other 50 (4.0%) | IPv4 1239 (99.5%) | IPv6 6 (0.5%) | In 849 (71.8%) | Out 334 (28.2%) | Deep Samples 1245 (100.0%)
Pkt Rates In 37/s 0/3/8/26 pps | Out 1/s 0/3/6/15 pps | IP Card. In: 54 | Out: 54

DNS Wire Pkts 302 (24.3%) | UDP 302 (100.0%) | TCP 0 (0.0%) | IPv4 296 (98.0%) | IPv6 6 (2.0%) | Query 157 (52.0%) | Response 145 (48.0%)
DNS Xacts 144 | In 49 (34.0%) | Out 95 (66.0%) | In 20.1/121.4/163.6/318.3 ms | Out 21.0/86.7/125.3/317.1 ms | Qname Card. 91
DNS NOERROR 145 (100.0%) | SRVFAIL 0 (0.0%) | NXDOMAIN 0 (0.0%) | REFUSED 0 (0.0%) | Time Window 6:43PM to 6:47PM, Period 258s

Top QName 2
.office.com                         36 (11.9%)
.google.com                         34 (11.3%)
.netflix.com                        28 ( 9.3%)
.spotify.com                                16
._tcp.local                                 13
.live.com                                   12
.googleapis.com                             12

Top QName 3
api-global.netflix.com              24 ( 7.9%)
.measure.office.com                 12 ( 4.0%)
dealer.spotify.com                  12 ( 4.0%)
.aria.microsoft.com                         10
textsecure-service.whispersystems.org       10
setup.icloud.com                             8
.ms-acdc.office.com                          8

Top NX

Slow In
api-global.netflix.com               1 ( 0.7%)

Top QTypes
A                                  268 (88.7%)
AAAA                                20 ( 6.6%)
PTR                                 11 ( 3.6%)
TXT                                          2

Top RCodes
NOERROR                           145 (100.0%)

Top SRVFAILS

Slow Out
edgeapi.slack.com                    1 ( 0.7%)
outlook.live.com                     1 ( 0.7%)
dealer.spotify.com                   1 ( 0.7%)
api-global.netflix.com                       1
outlook.office.com                           1

Top REFUSED

IPv4
127.0.0.1                          519 (41.7%)
192.168.0.189                      374 (30.0%)
192.168.0.114                       98 ( 7.9%)
192.168.0.55                                62
216.239.32.10                               12
216.239.36.10                                8
216.239.34.10                                6

IPv6

Top DNS UDP Ports
5353                                14 ( 4.6%)
48279                                2 ( 0.7%)
21952                                2 ( 0.7%)
65101                                        2
33225                                        2
50842                                        2
12213                                        2

Top GeoLoc
Unknown                           1055 (84.7%)
NA/United States                    78 ( 6.3%)
NA/United States/CA/Mountain View 30 ( 2.4%
NA/United States/WA/Redmond                  6
NA/United States/CA/Los Angeles              4
NA/United States/CA/Sacramento               2
EU/Ireland/L/Dublin                          2

Top ASN
Unknown                           1055 (84.7%)
15169/GOOGLE                        30 ( 2.4%)
8068/MICROSOFT-CORP-MSN-AS-BLOCK 26 ( 2.1%)
8075/MICROSOFT-CORP-MSN-AS-BLOCK            24
16509/AMAZON-02                             14
15133/EDGECAST                               8
21342/Akamai International B.V.              6
```

Command Line UI

Centralized Operation

Centralized Collection

Generic telegraf HTTP input plugin

⇩

Elasticsearch output plugin

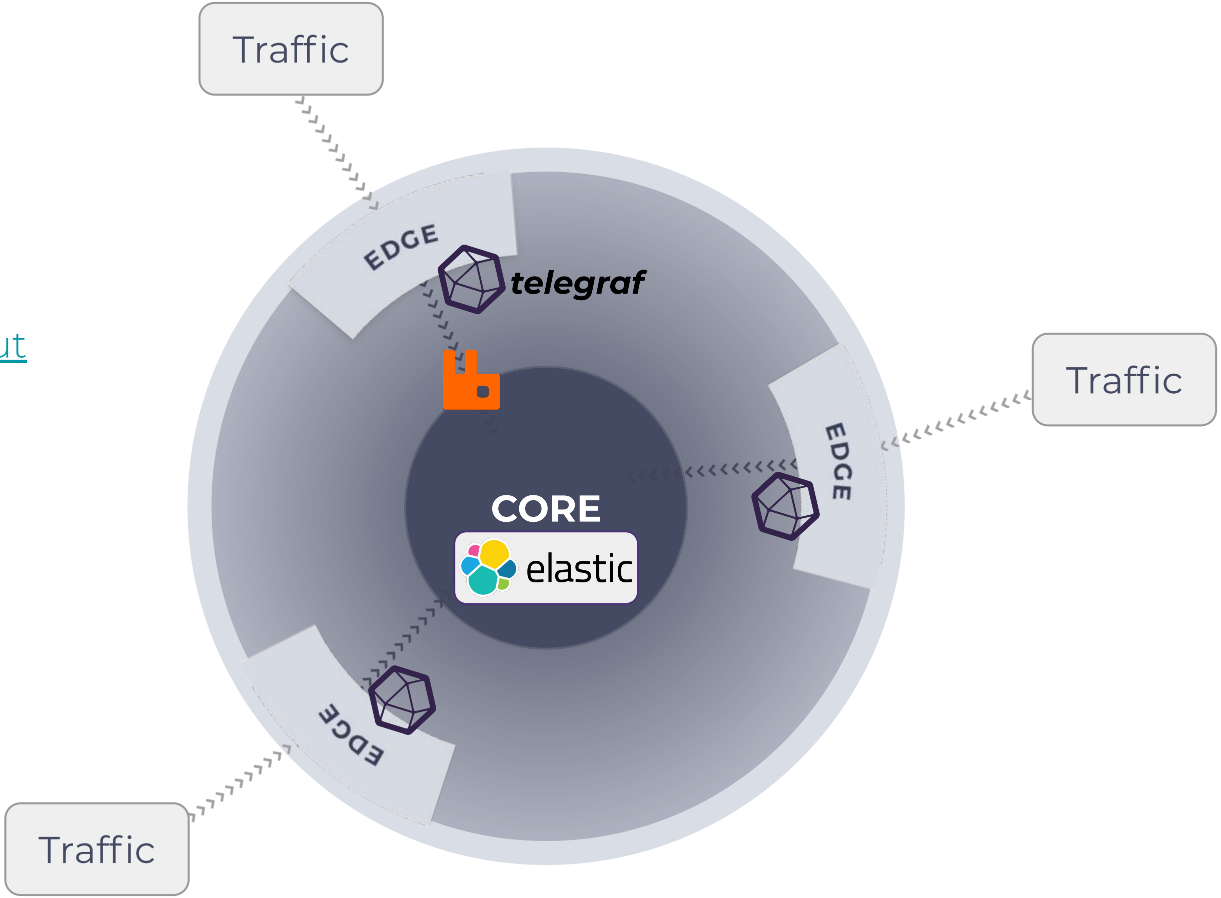

NS1.

Grafana Dashboard

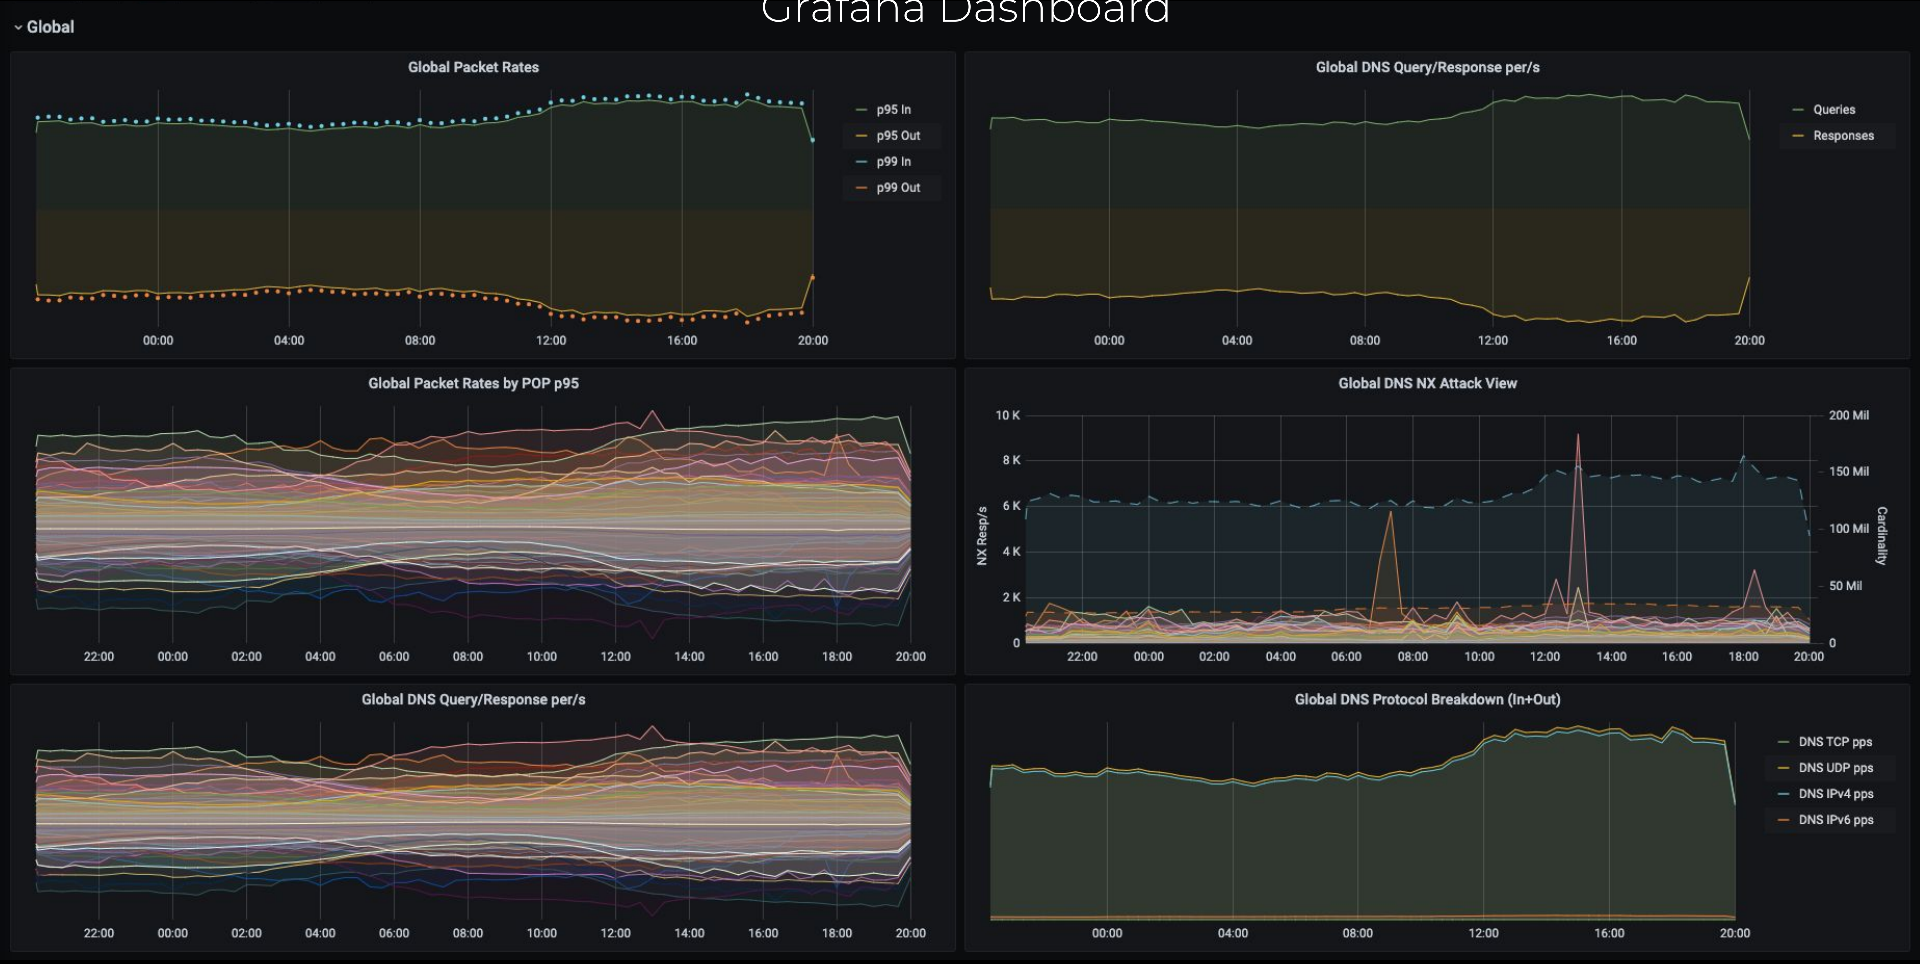

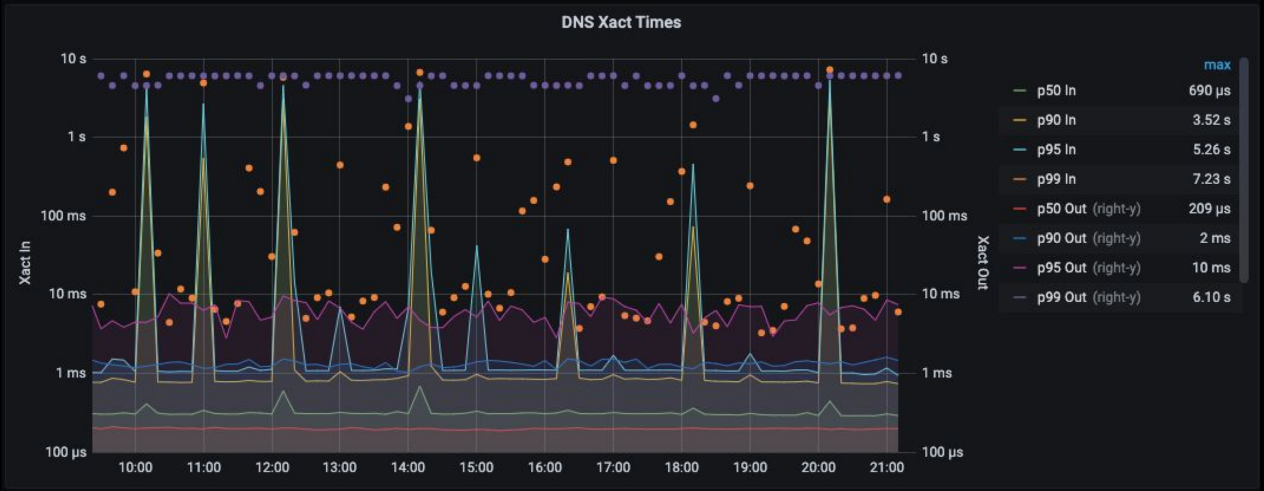

DNS Xact Times

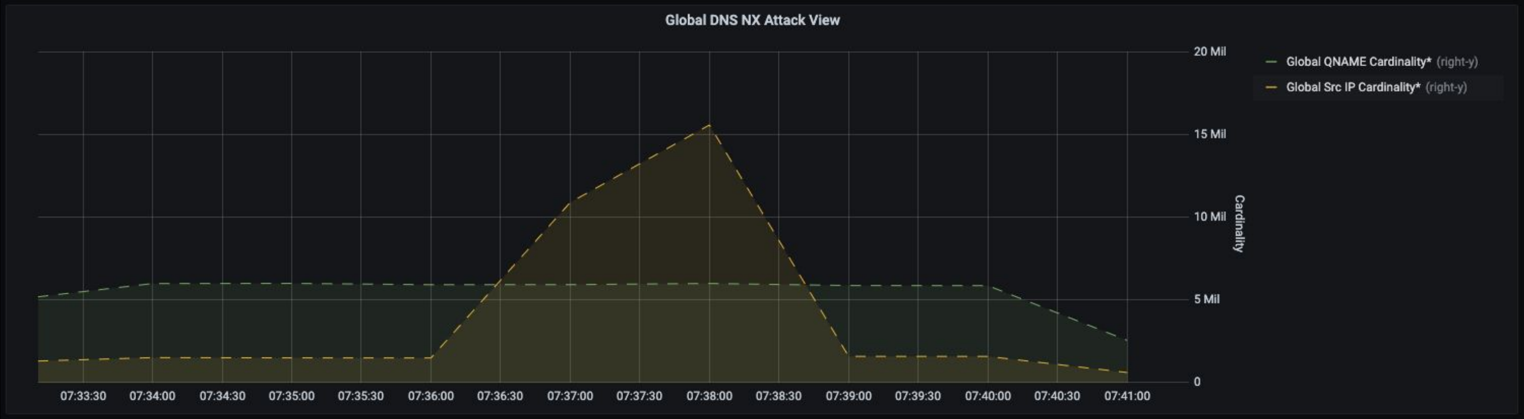

Global DNS NX Attack View

Centralized Top-N

- Requires **map reduce**
- Visualization not supported in Grafana or even Kibana
- Map reduce script for Elasticsearch supplied on Github
- **Calculates global Top N for all tables** across selected nodes and time frame
- Stand alone dashboard to visualize this data in progress
    - based on bokeh

Open Source

NS1.

GitHub

github.com/ns1/pktvisor

Collaborators welcome!

NS1.

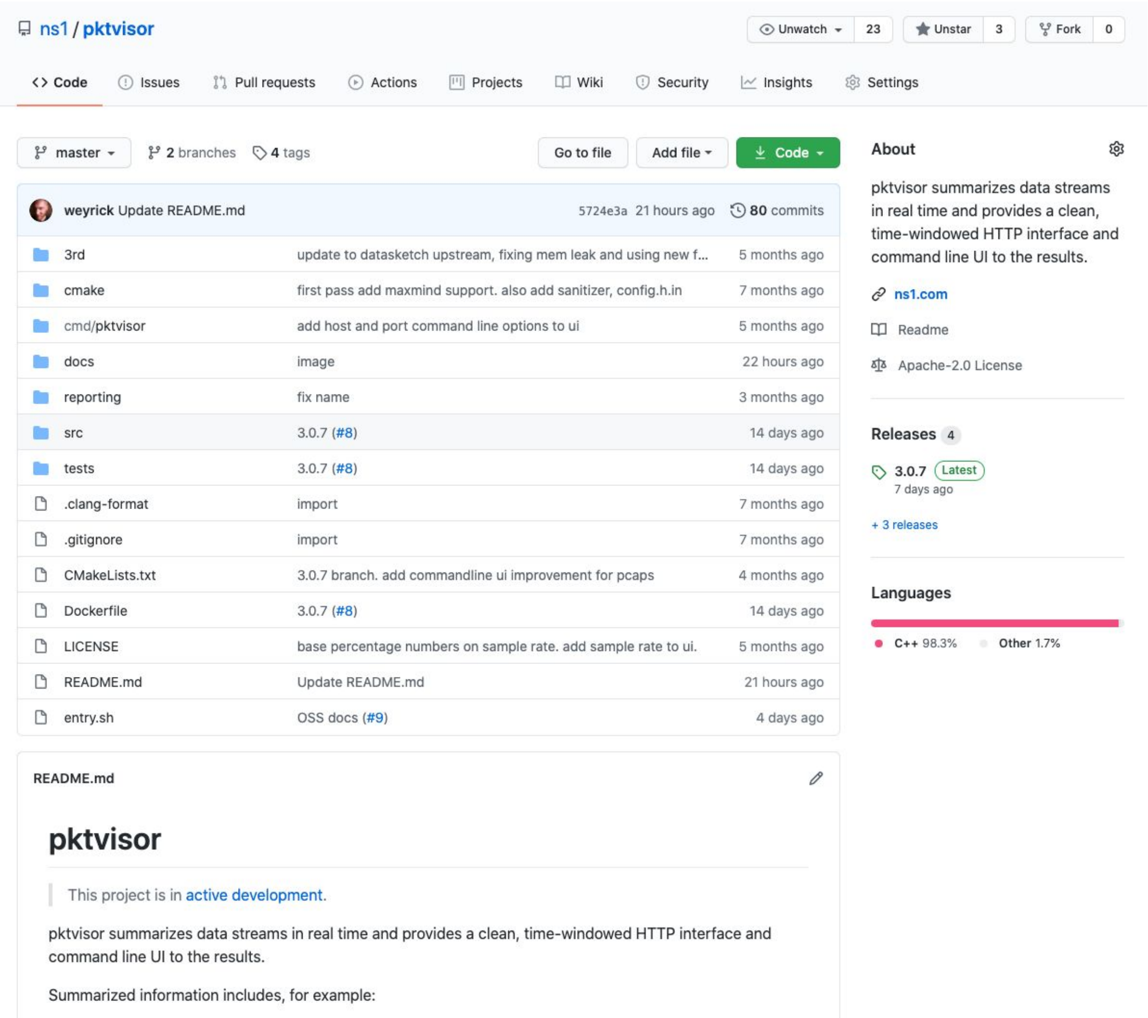

weyrick Update README.md 5724e3a 21 hours ago 80 commits

3rd	update to datasketch upstream, fixing mem leak and using new f...	5 months ago
cmake	first pass add maxmind support. also add sanitizer, config.h.in	7 months ago
cmd/pktvisor	add host and port command line options to ui	5 months ago
docs	image	22 hours ago
reporting	fix name	3 months ago
src	3.0.7 (#8)	14 days ago
tests	3.0.7 (#8)	14 days ago
.clang-format	import	7 months ago
.gitignore	import	7 months ago
CMakeLists.txt	3.0.7 branch. add commandline ui improvement for pcaps	4 months ago
Dockerfile	3.0.7 (#8)	14 days ago
LICENSE	base percentage numbers on sample rate. add sample rate to ui.	5 months ago
README.md	Update README.md	21 hours ago
entry.sh	OSS docs (#9)	4 days ago

README.md

pktvisor

> This project is in active development.

pktvisor summarizes data streams in real time and provides a clean, time-windowed HTTP interface and command line UI to the results.

Summarized information includes, for example:

About

pktvisor summarizes data streams in real time and provides a clean, time-windowed HTTP interface and command line UI to the results.

ns1.com

Readme

Apache-2.0 License

Releases 4

3.0.7 Latest
7 days ago

+ 3 releases

Languages

● C++ 98.3% ● Other 1.7%

SMARTBEAR SwaggerHub™

app.swaggerhub.com/apis/ns1labs/pktvisor

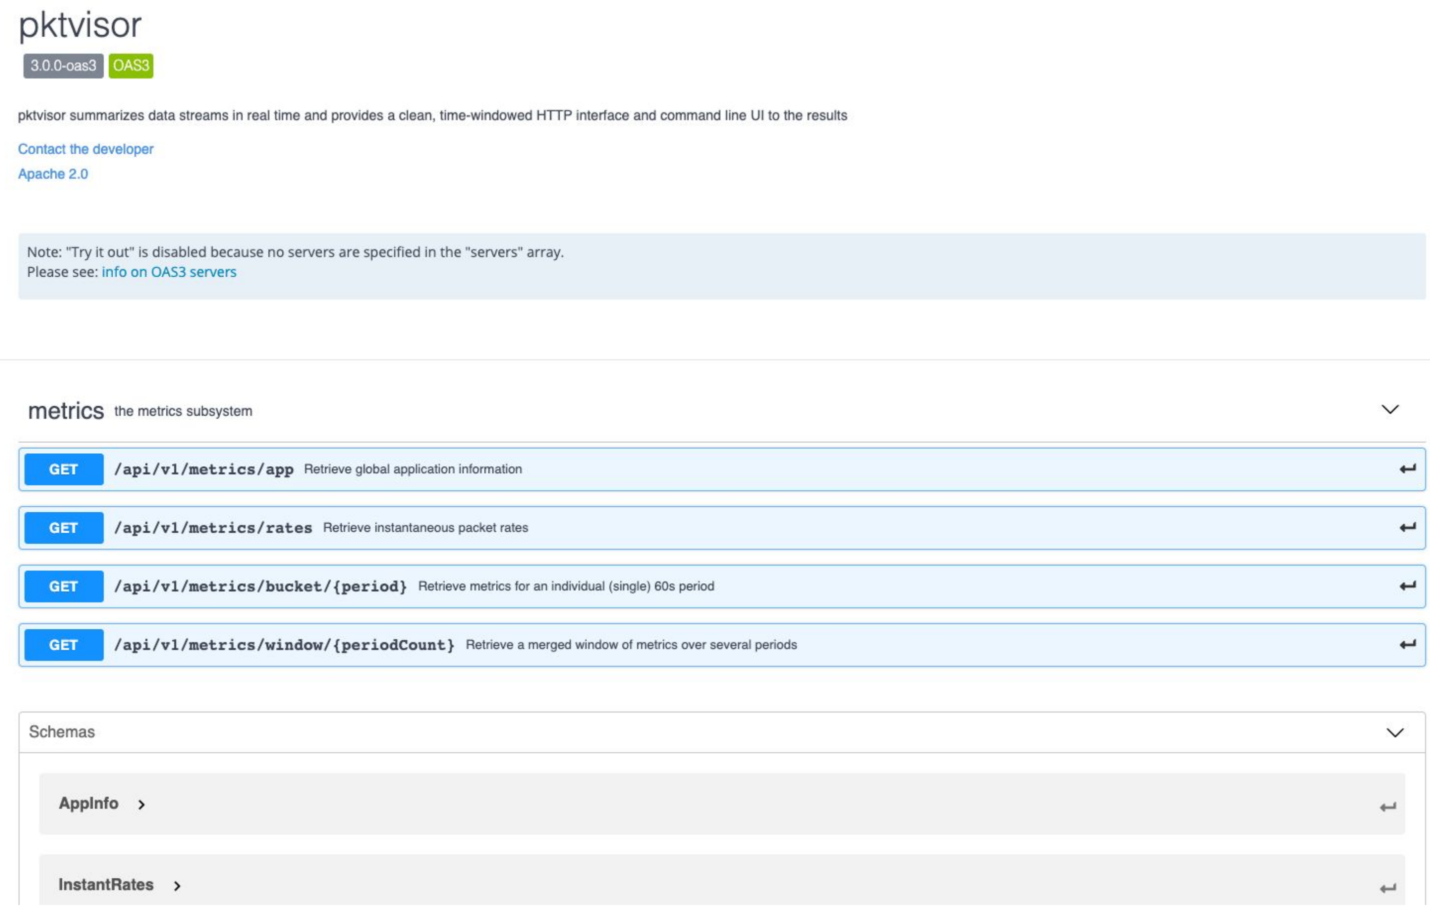

pktvisor

3.0.0-oas3 OAS3

pktvisor summarizes data streams in real time and provides a clean, time-windowed HTTP interface and command line UI to the results

Contact the developer

Apache 2.0

Note: "Try it out" is disabled because no servers are specified in the "servers" array.
Please see: info on OAS3 servers

metrics the metrics subsystem

- GET `/api/v1/metrics/app` Retrieve global application information
- GET `/api/v1/metrics/rates` Retrieve instantaneous packet rates
- GET `/api/v1/metrics/bucket/{period}` Retrieve metrics for an individual (single) 60s period
- GET `/api/v1/metrics/window/{periodCount}` Retrieve a merged window of metrics over several periods

Schemas

- AppInfo ›
- InstantRates ›

Future plans and ideas

- Central control plane with dynamic policies
  - Current REST API is read only, but with writes we enable dynamic control
- Allow other types of streams
  - it can summarize data sent from e.g. a message queue or pipe
- Allow summarizing other protocols
  - e.g. DoH, DHCP, AMQP, ...
- Novel metrics
  - eBPF probe module for summarizing application information
- Optimization, targeting 100’s k QPS with dynamic sampling rate

Thank You!

Questions?

Shannon Weyrick, VP Architecture @ NS1
sweyrick@ns1.com