On The Edge of Small Data

Shannon Weyrick
VP Research/Fellow · NS1
sweyrick@ns1.com

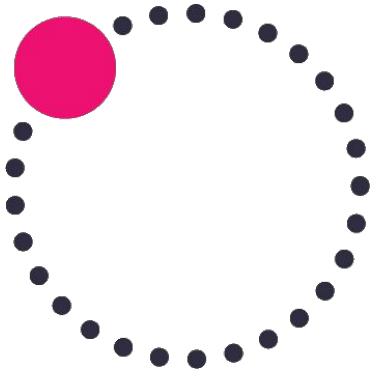

NS1.

preface: the case for small data

NS1.

NS1 Case Study

- Managed Authoritative DNS with 26 Global Anycasted POPs
- **>100 billion** DNS queries per average day
- **29 TB** of raw DNS wire packets per day

https://orb.community

The Data Conundrum

What we think we want:
All The Data

...because we think we *may* use it all *someday*

What we actually want:
Targeted Insights

...to help us operate, debug, scale and protect our networks *today*

There is a price to pay for streaming raw data to a central solution

The Costs of Raw Data

- Complicated data pipelines for centralized collection
- Batch processing costs to make it actionable
- Inability to make sense of or take advantage of all the data
- Slow dashboards, short retention times
- Slow reaction times to critical events
- Ingestion costs (esp. SaaS)

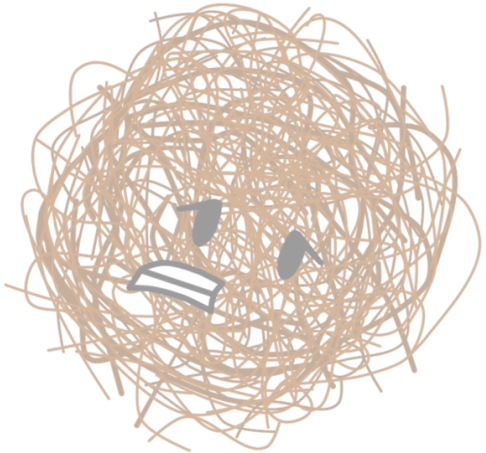

Paradigm Shift: Small Data

- Push the conversion of raw → actionable to the edge
  - Distribute as close to the source as possible
- React quicker
  - Make those insights available at the edge *and* centrally
- Collect, process and store less
- Don’t find the needles in the haystack: just collect the needles
- Dynamically decide what your team needs at any time

Shannon Weyrick

Orb Founder, VP Research @ NS1

- 26 years in industry, 8 years at NS1
- NS1 engineering leadership
- Since start of 2021 focused on Orb open source innovation @ NS1 Labs
- sweyrick@ns1.com

NS1.

https://orb.community

If you remember just one thing from this talk...

NS1.

Orb is Open Source Edge Observability

- **Observability tool** designed for **distributed edge networks**
- Uses **small data** paradigm with **dynamic policy orchestration**
- Real-time **insights** into **data flow** on the **distributed edge**
- Integrates with **modern observability stacks**
- **Free** and **open source**, backed by NS1

Last 2 days | 1m

Agents | blesk + gw | Policy | general | Filters | +

pktvisor docs

DNS Overview (Now)

Agents	Total QPS (p95)	Xact In	Xact In Time	Xact Out	Xact Out Time	Unique QNames
Shown 2 / Total 2	6	22	p95 53.1 ms / p99 238 ms	17	p95 50.4 ms / p99 50.4 ms	19

DNS Name Overview

Top DNS Names

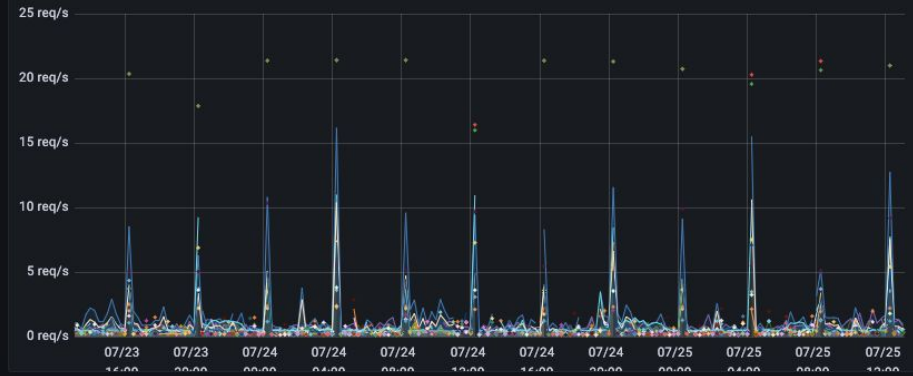

	avg
.nflxvideo.net	17.5 req/s
.oca.nflxvideo.net	11.2 req/s
.icloud.com	6.26 req/s
.push.apple.com	3.59 req/s
.skype.com	3.35 req/s
push.apple.com	3.26 req/s
.whatsapp.net	2.17 req/s
.trafficmanager.net	1.86 req/s
.akadns.net	1.72 req/s
.netflix.com	1.65 req/s
.gvt1.com	1.56 req/s

Top NXDOMAIN

Metric	Value (sum)
iad07.nsone.co	32
6b7b001b-ec4a-fbd...	13
epc.geo.mnc260.m...	11
lb._dns-sd._udp.mo...	11
lb._dns-sd._udp.0.1....	5
apple-cloudkit.fe.ap...	5
local	4
lb._dns-sd._udp.0.0....	2

Top SRVFAIL

Metric	Value (sum)
api.fitbit.com	
epdg.epc.mnc240....	

DNS QName Tables

Names Agg3

Metric	Value (sum)
.logs.roku.com	4.88 k
.com.akadns.net	3.35 k
.elb.amazonaws.com	2.09 k
com.akadns.net	1.94 k
.fe.apple-dns.net	1.82 k
.clients6.google.com	1.72 k

Top REFUSED

Metric	Value (sum)
unity.rokuapi.net	

Result Codes

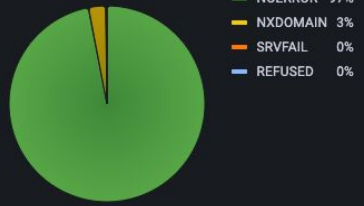

QTypes

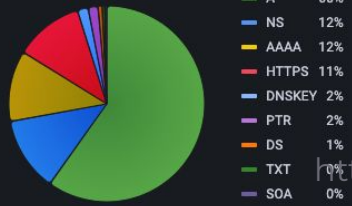

Transport

Deep Streaming Analysis

sample of current metrics (out of 110 + growing)

Network (L2-L3)

- Top IPs
- Top Geo/ASN
- IP Cardinality
- Rate Perctil.
- Payload Perc.
- Throughput
- Protocol
- ...

DNS

- Top QNames
- Geo by ECS
- Transactions
- Protocol
- Rate Perctil.
- QName Card.
- Timer Perctil.
- Amplification
- ...

Flow

- Top flows
- Flow rates
- Protocols
- Ports
- Interfaces
- Throughputs
- ...

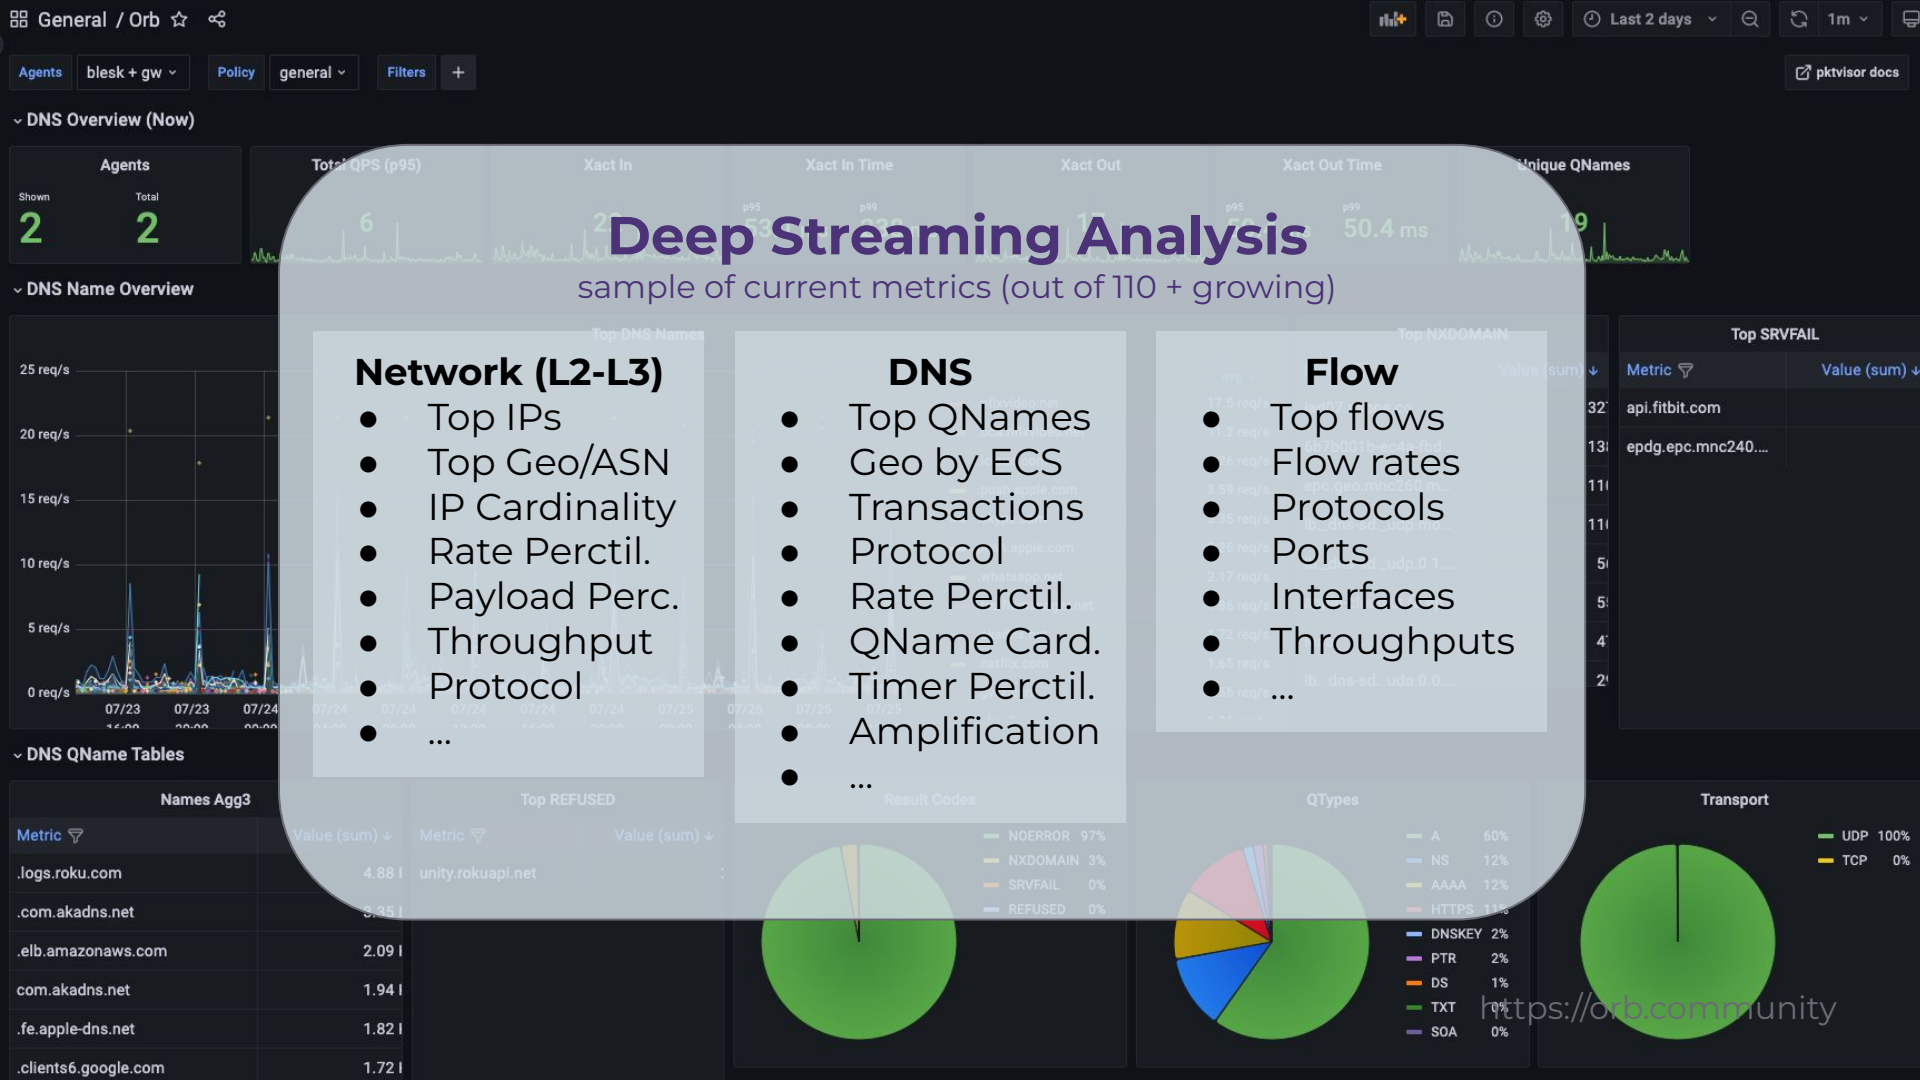

control tower for the edge

Orb control plane: cloud native application

NS1.

Control Tower for Dynamic Edge Observability

- **Usability & Automation**: Portal UI & REST API
- **Fleet management**: connect, organize, and manage edge agents
- **Policy management**: recipes for analyzing data streams, deciding which agents should receive which policies in real time
- **Data collection & Sinking**: scrape lightweight metric output from all policies across all agents and push to the proper databases and dashboards

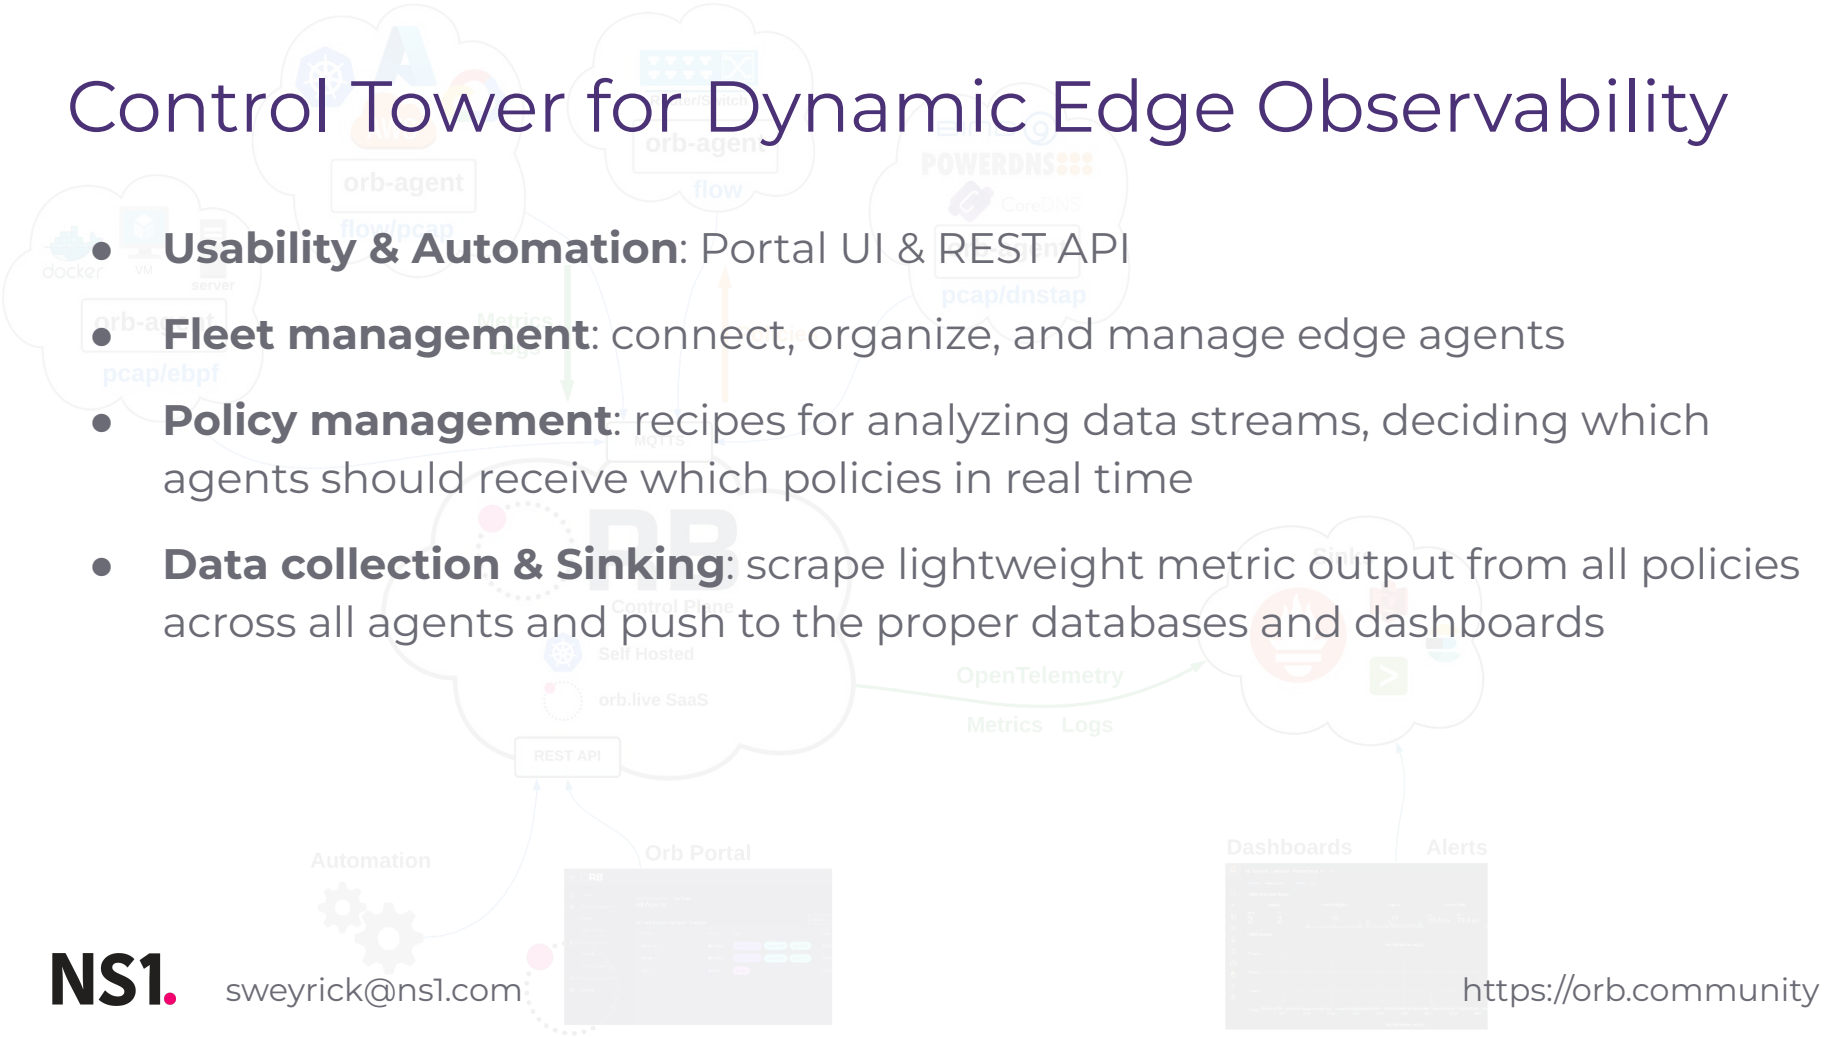

Orb Architecture Overview

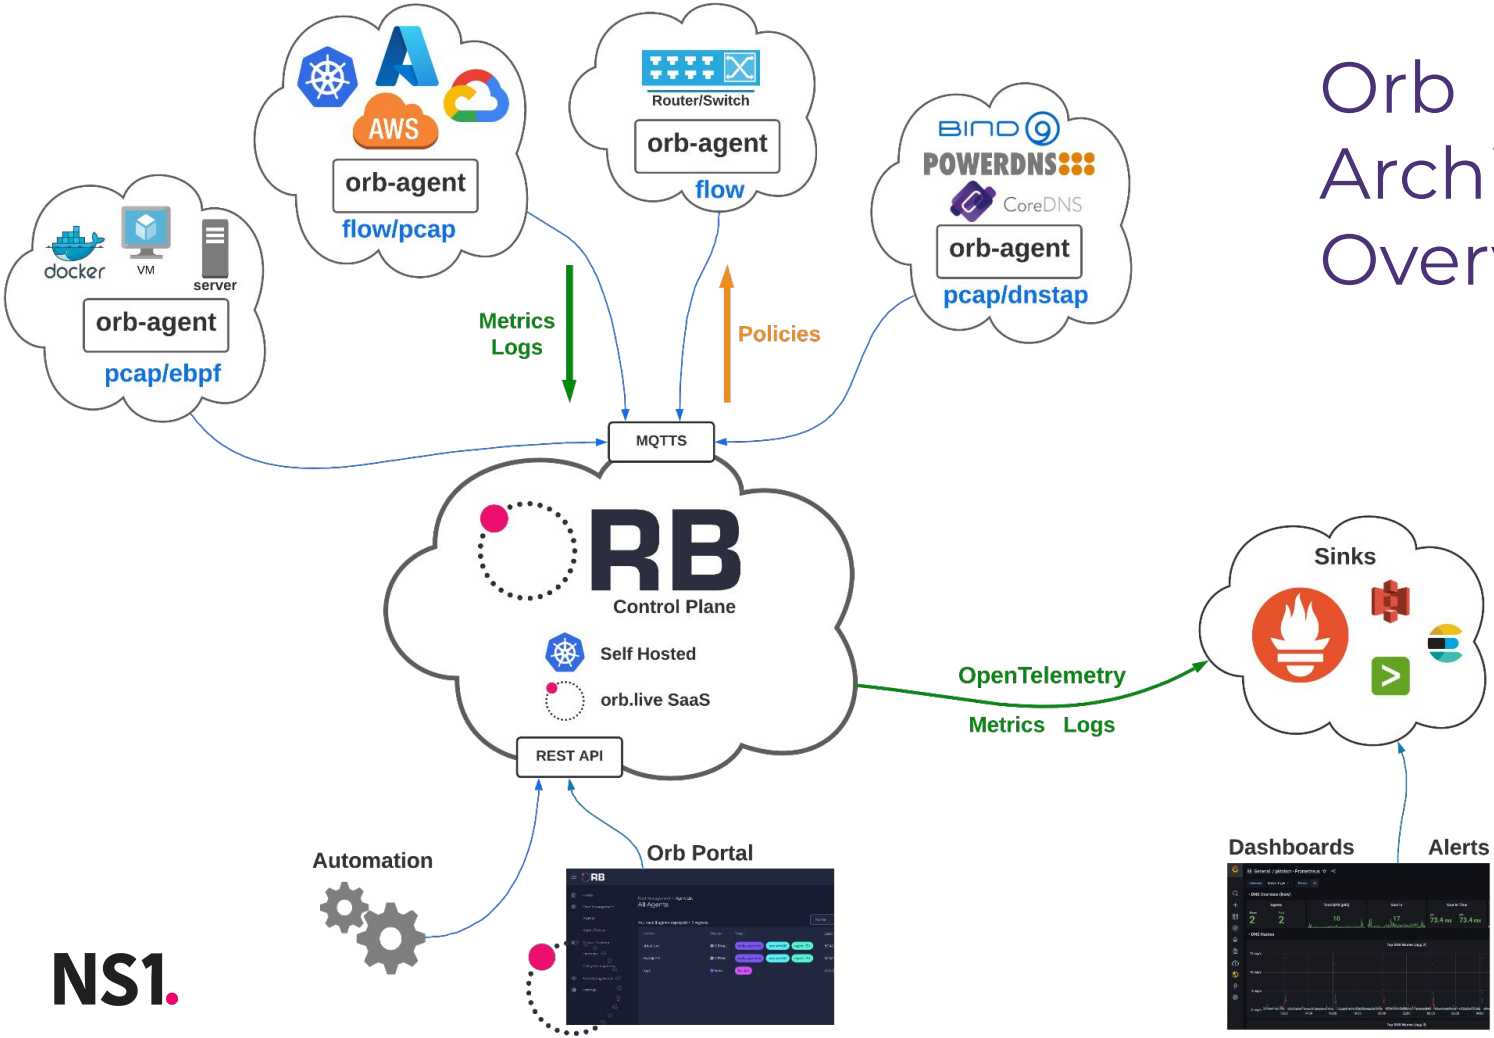

NS1.

https://orb.community

Fleet Management

Connect, organize, and manage edge agents

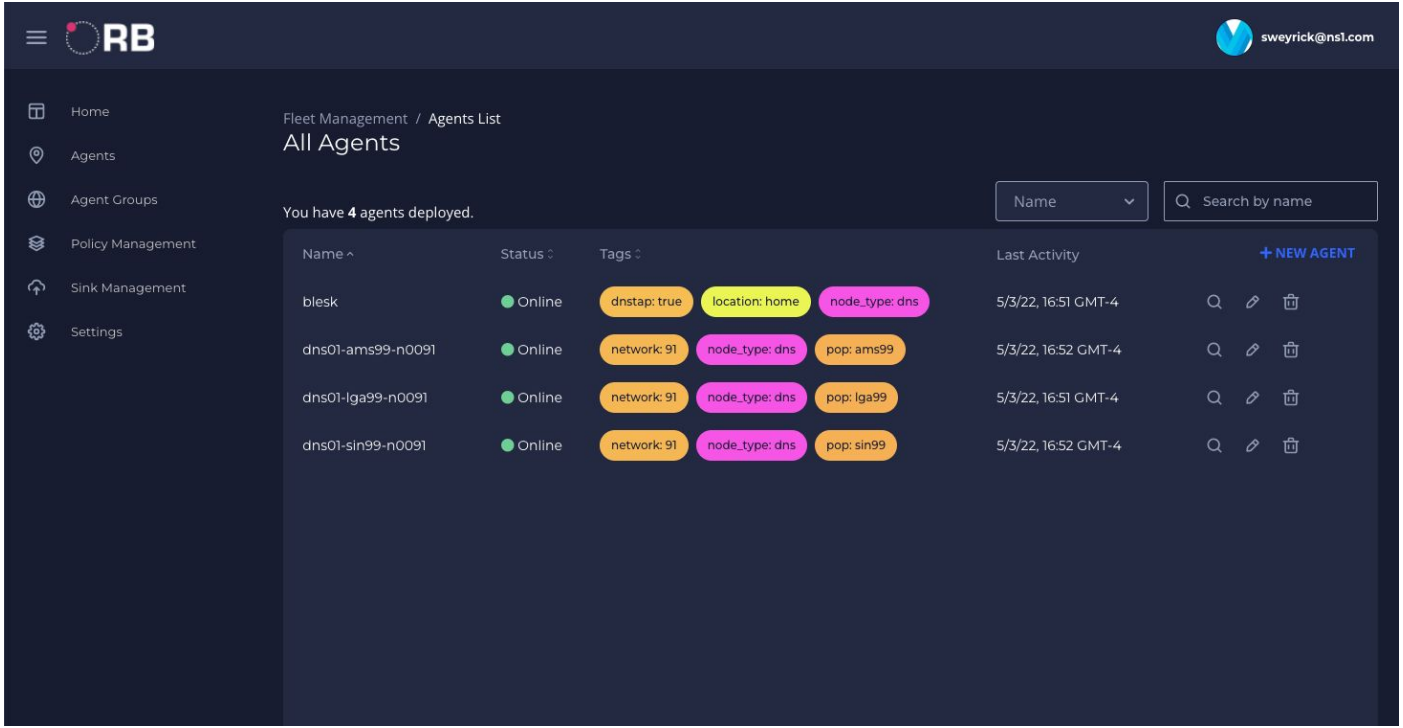

Policy Management

Recipes for analyzing data streams

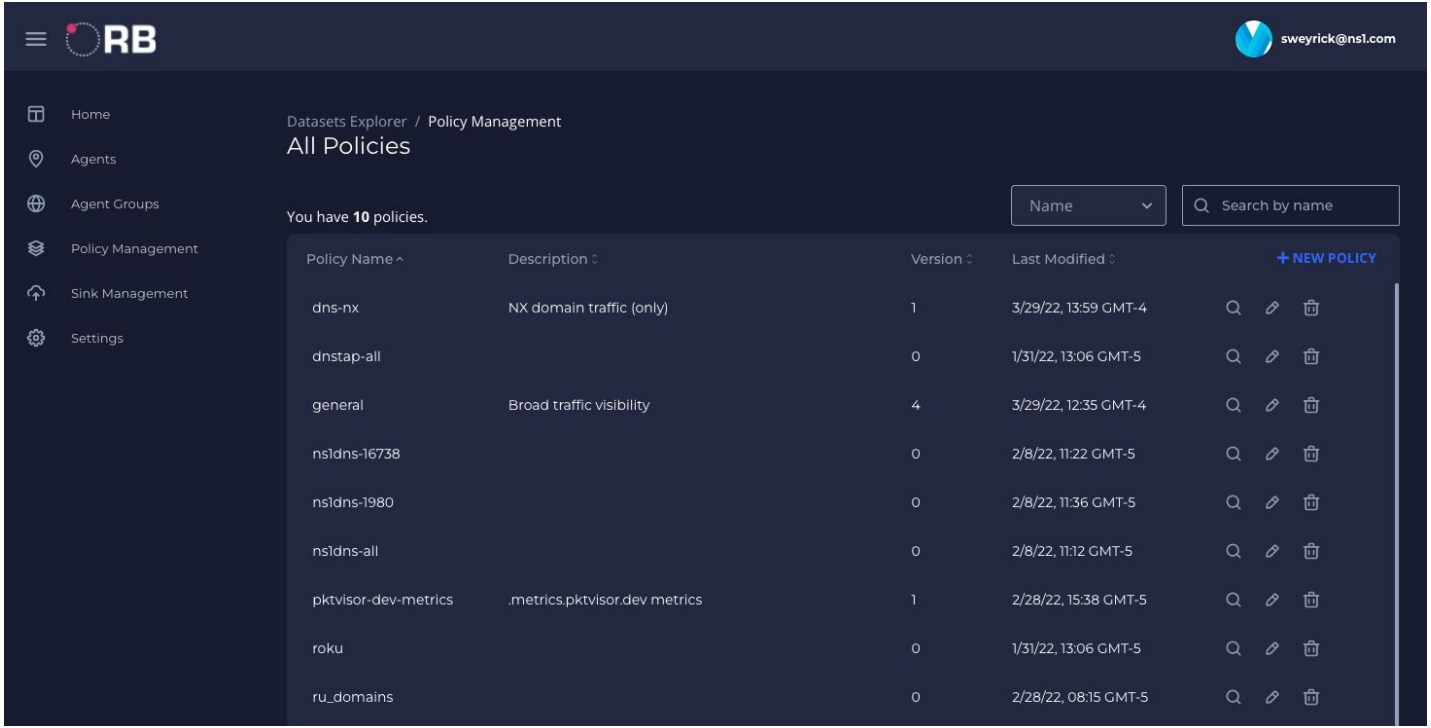

Sink Management

Which databases and dashboards to send metrics to

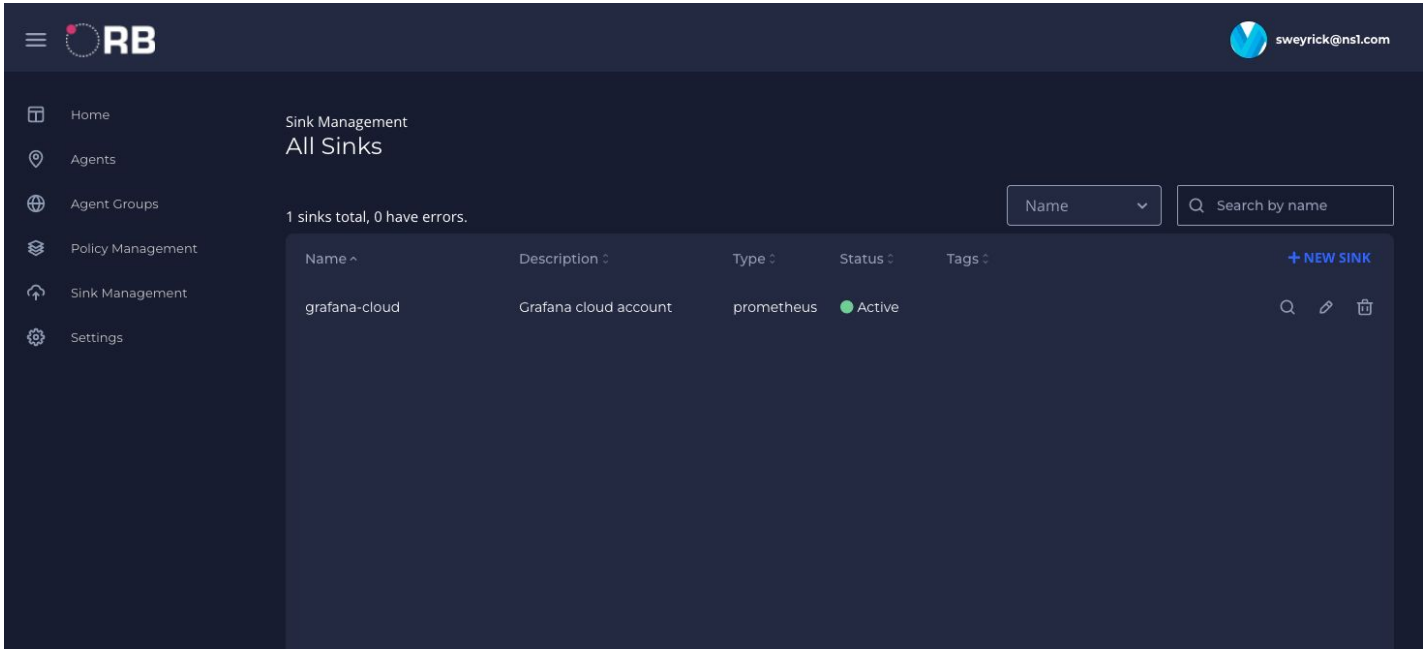

Configuration Management

Which agents should run which policies, update in real time

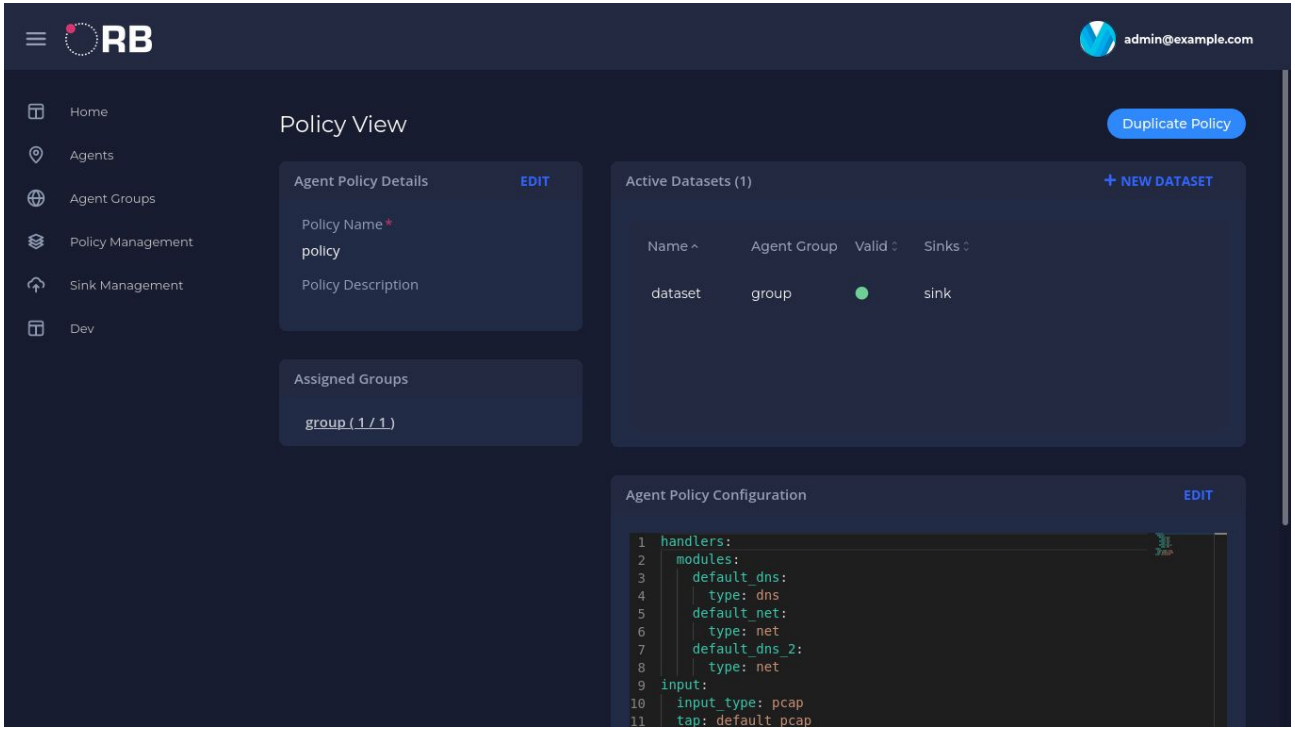

Data Collection & Sinking

Scrape lightweight metric output from all policies across all agents and push to the proper databases and dashboards

NS1. sweyrick@ns1.com

https://orb.community

edge agent for streaming analysis

orb-agent

NS1.

What Is The Orb Edge Agent?

- **Taps into** multiple, concurrent data streams at the edge
- Uses **fast streaming algorithms** to **analyze deeply** in real time
- **Efficiently summarizes** important insights, generate metrics
- Can be **reprogrammed in real time** with dynamic policies
- Can **scale up** and **scale down**

What Can It Tap Into?

- Packet capture (with BPF filters)
- dnstap (socket, TCP)
- Network flow (sFlow, Netflow/IPFIX)
- future: envoy taps, eBPF, WebAssembly...
- Expandable via custom loadable modules

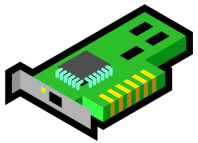

sFlow

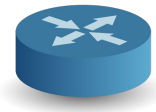

DNS

What Can It Generate Metrics For?

- L2-L3 Network
- DNS
- DHCP
- Flows
- Policy resource usage
- Expandable via custom loadable modules

Orb Edge Agent

Embedded Stream Processing

powered by pktvisor

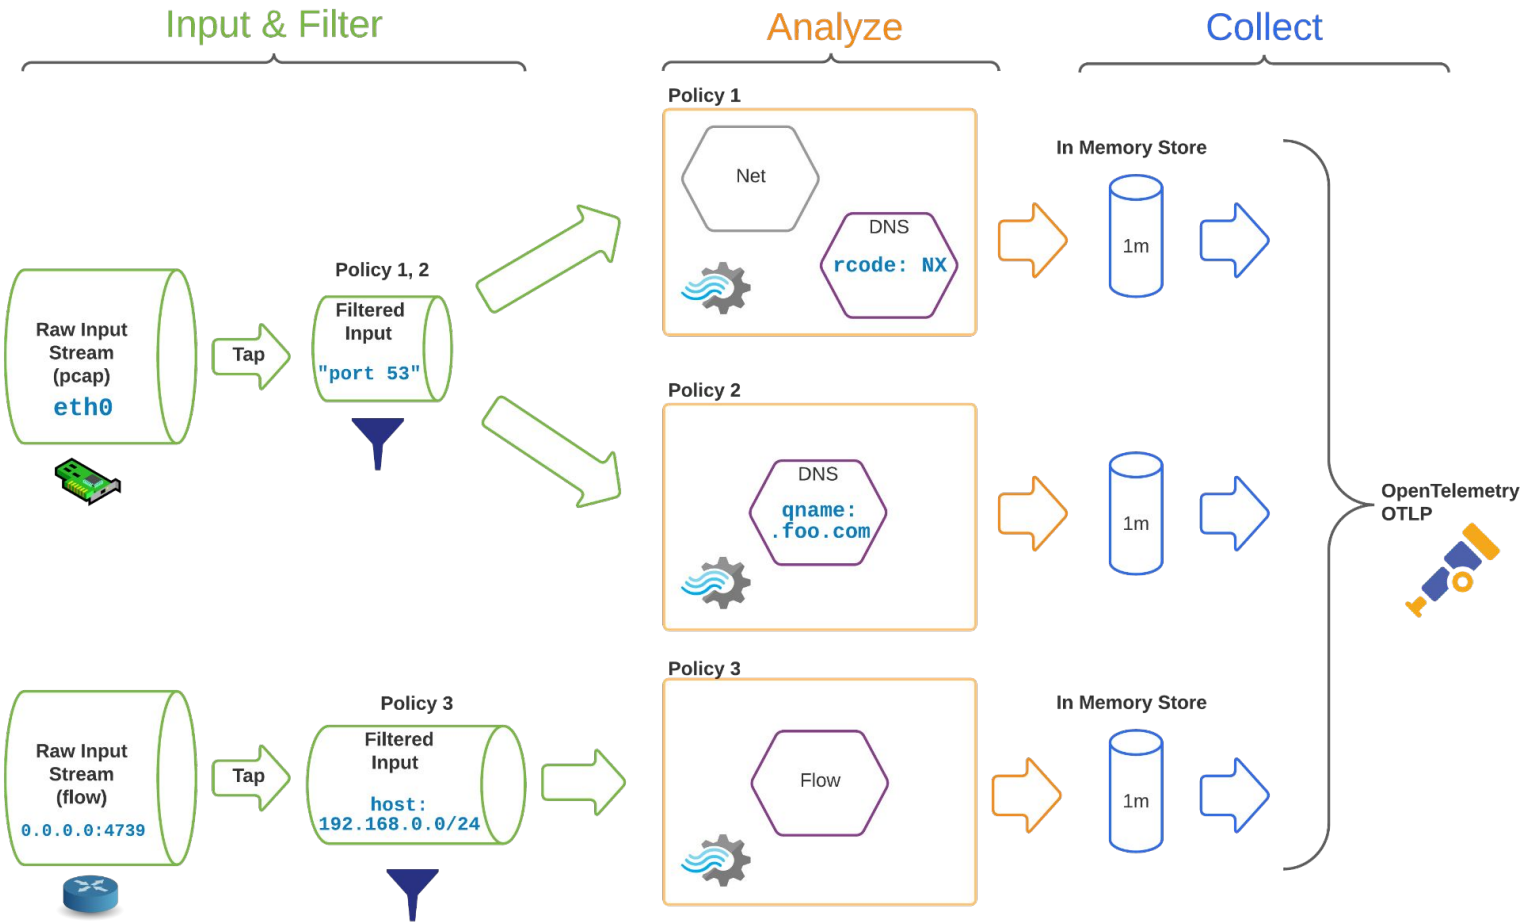

NS1. sweyrick@ns1.com

https://orb.community

Use Case: DNS Analysis

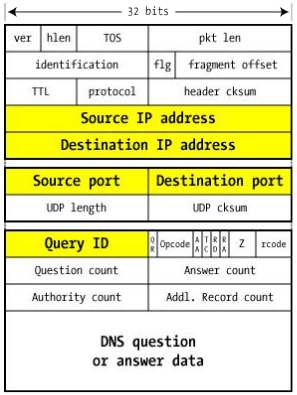

DNS packet on the wire

raw DNS traffic

pktvisor

Streaming Analysis

```
Top Queries
Top Query Types
Top Result Codes
Rate Percentiles
Top Sources
(according to policy)
...
```

1 minute summaries

flow

Use Case: DNS Analysis

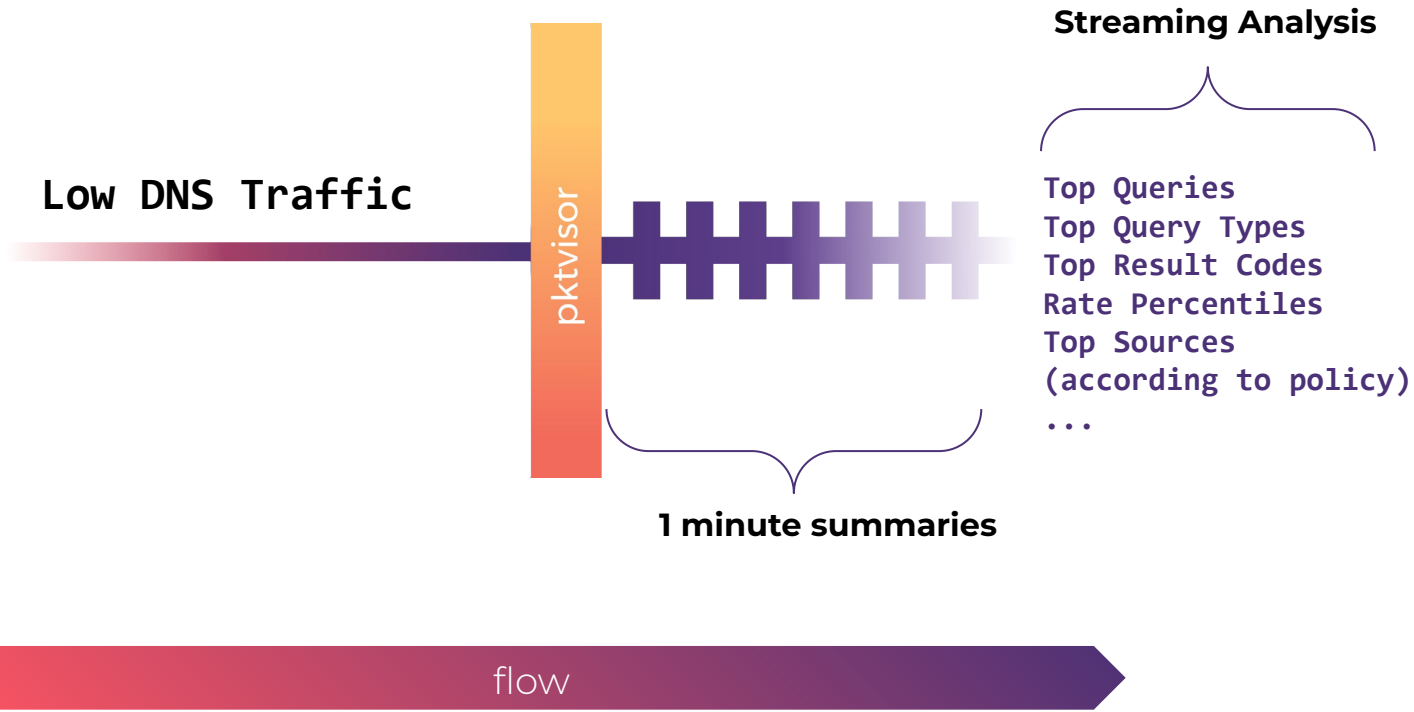

Use Case: DNS Analysis

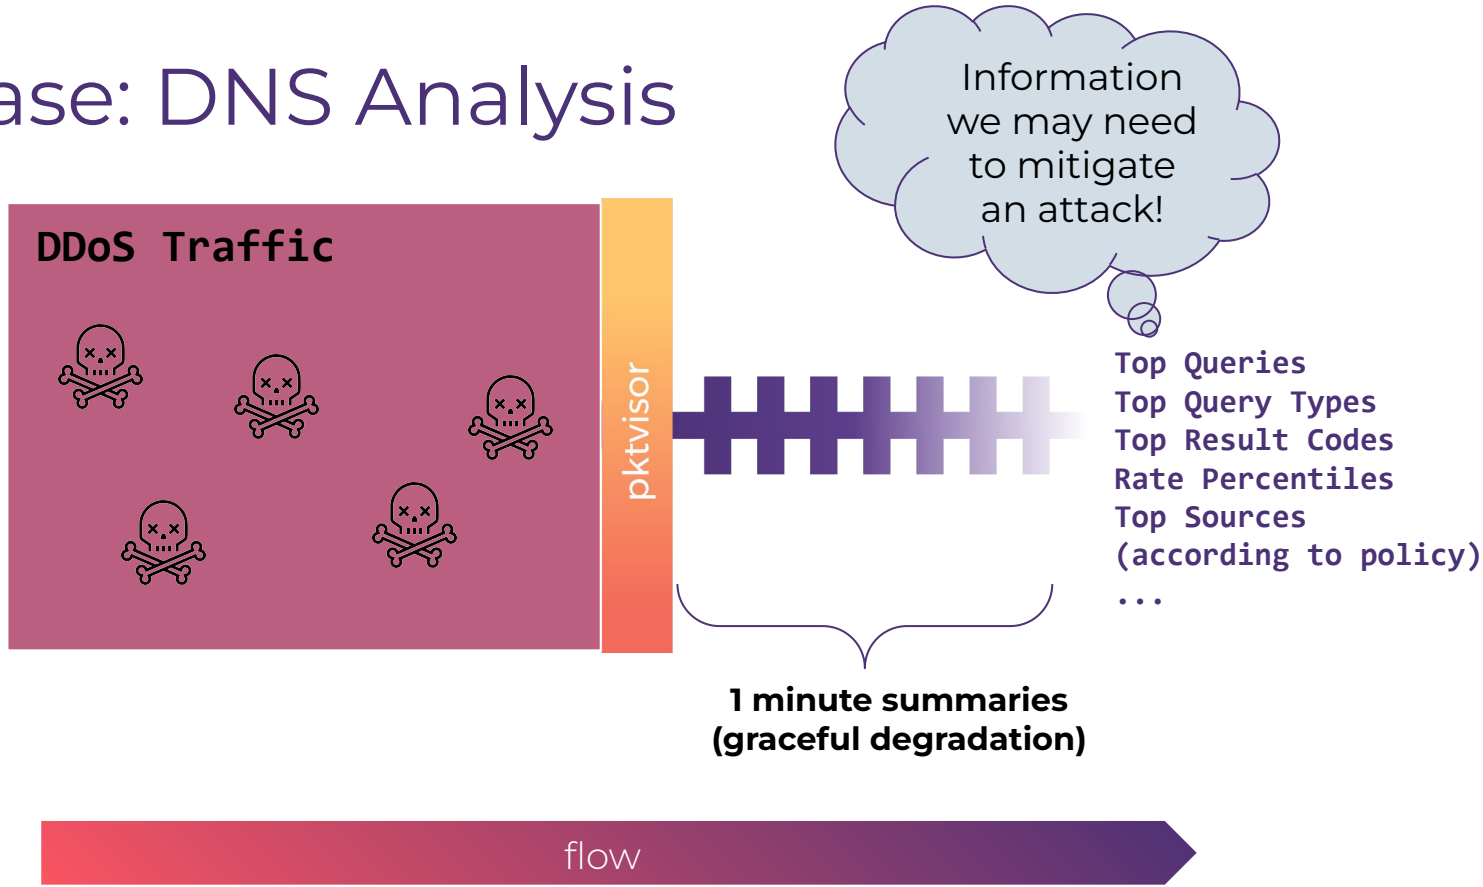

project status

NS1.

Open Source, Backed by NS1

- Core tech has **powered NS1 operational** visibility for years
- Powers NS1 **DNS Insights** product for NS1 customers today
- All development done on **GitHub**
- Now working with **design partners** to develop the project towards various edge network, IT, and NetDevOps use cases
- **Join us!**

Exciting Future

- Expanding our active **community** and **contributions**
- New input stream sources and **analyzers**
- **Machine learning** for automated insights and anomaly detection
- **pcap samples** from distributed fleet orchestrated from control plane
- Policy driven **actions on the edge**
- What are your ideas? Let us know!

conclusion

NS1.

Remember This

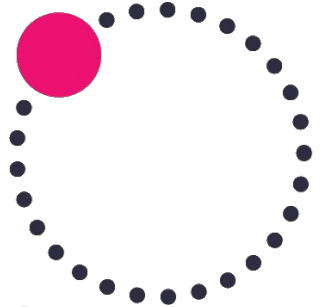

- **Observability tool** designed for **distributed edge networks**
- Uses **small data** paradigm with **dynamic policy orchestration**
- Real-time **insights** into **data flow** on the **distributed edge**
- Integrates with **modern observability stacks**
- **Free** and **open source**, backed by NS1
- You can **POC in 10 minutes** on orb.live

Do This

- Join the community: https://orb.community
- Try Orb SaaS for free: https://orb.live
- Star the project: github.com/ns1labs/orb
- Contact us! We're seeking development partners like you.

thank you

NS1.