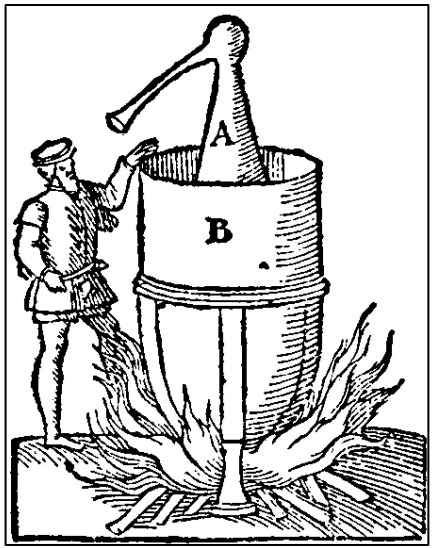

High-Proof Data

TSDB Distillation of Recursive DNS Metrics

Alexis Fasquel
alexis@pch.net

Introduction

Our networks

- Authoritative Resolvers
- Public Recursors (pilot)

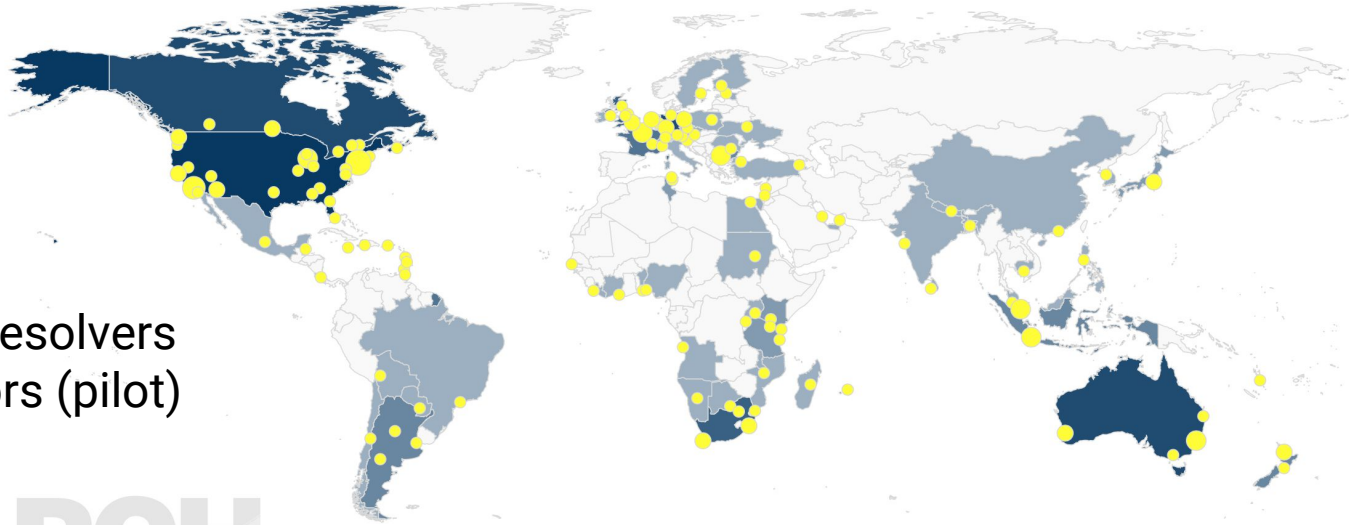

154 POPs around the world

Introduction

PCH running pilot for anycast public DNS recursor

- Improving DNS security (optional RPZ/blocklist from >20 threat providers, DNSSEC, etc.)
- Provide user privacy (PII never transmitted/sold, not a commercial project with profit motives, encryption support)
- Performance (large footprint, “close” to most Internet users)

Introduction

DNS Traffic

Monitoring needs

- Understand the routed nature of queries coming (transit, origin)
- Keep track of unique DNS records for easier malware domains detection
- Domain name monitoring (experimental operational req's)

The challenges

- **Scalability**
  - Many resolver clusters
  - Lots of queries - how to process without bottlenecks?
- **Resources**
  - Hardware/network limitations
  - Human resources and cost
  - Structured databases not well suited for this scale
- **Encryption**
  - PCAP monitoring of DNS becomes obsolete (we hope!)
- **PII Constraints**
  - No client IP addresses stored to disk or transmitted outside POP

The solution

Answer pre-defined questions at the edge

- **Specific uses cases not necessarily easy to solve**
  - Difficulties in joining BGP and DNS data
  - Look at the entire feed of data
- **Data retention**
  - Avoid an overgrown logging system
- **Allow distillation of data at the edges**
  - Fewer resources needed at core
  - Less flexibility for future questions

How to store data?

TSDB

- **Great for live monitoring**
  - Simple querying and easy visualisation
  - Simple data retention
  - Tagging becomes very important on the front end
- **Relatively easy to maintain**
  - Picking important metrics is critical
  - Scaling issues with long-term storage (data must be averaged)

How to move data?

Streaming protobuf!

- **DNSDist**
  - DNS load balancer
  - Stream DNS queries through protocol buffers
- **Protocol buffer duplication**
  - Multiple applications consuming the DNS pbuf stream
  - Custom middleware

Architecture

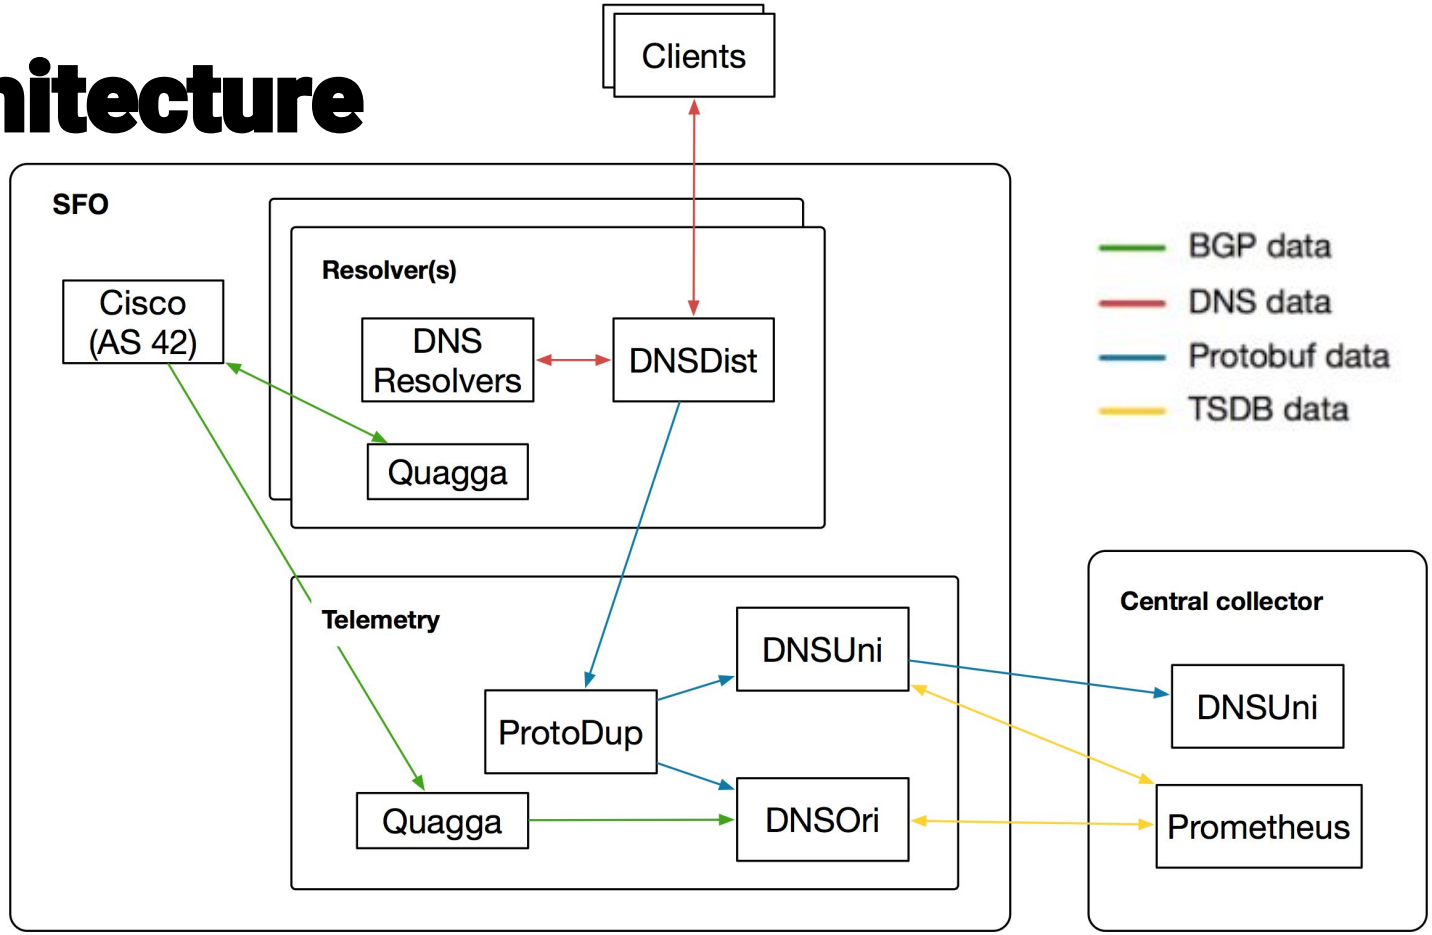

Results and current progress

DNSOri

- **Resolving BGP data in each POP**
  - Greatly distills the amount of data to move around
  - Manual de-aggregation if desired
- **High cardinality!**
  - Turns billions of points into low thousands series
  - Still hard on TSDBs - might need more than 1 central instance
- **Questions this data answers:**
  - Who should we peer with?
  - What are our changes in traffic?
  - Are we seeing split origin traffic? How often?
  - What auths send us the most garbage?

POP pao | Env prod | Host All | TopN 10

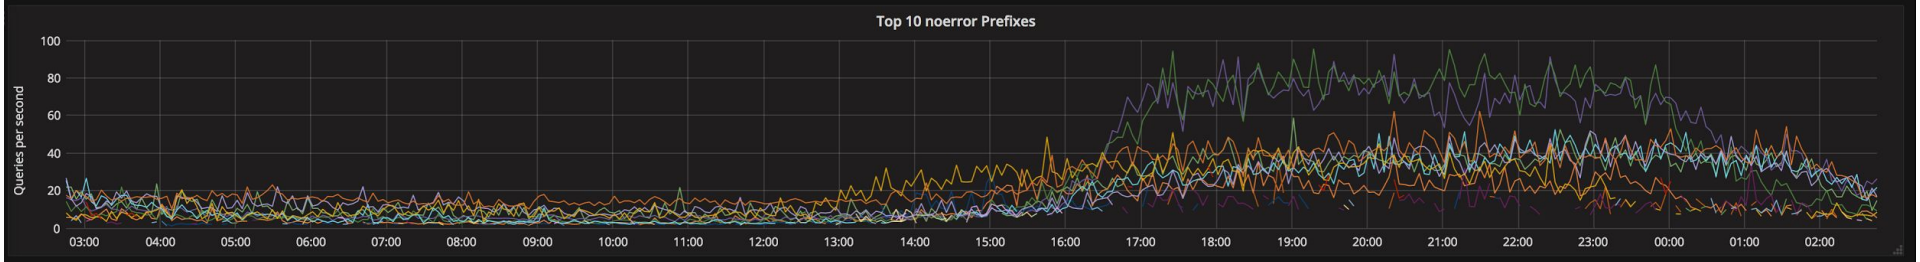

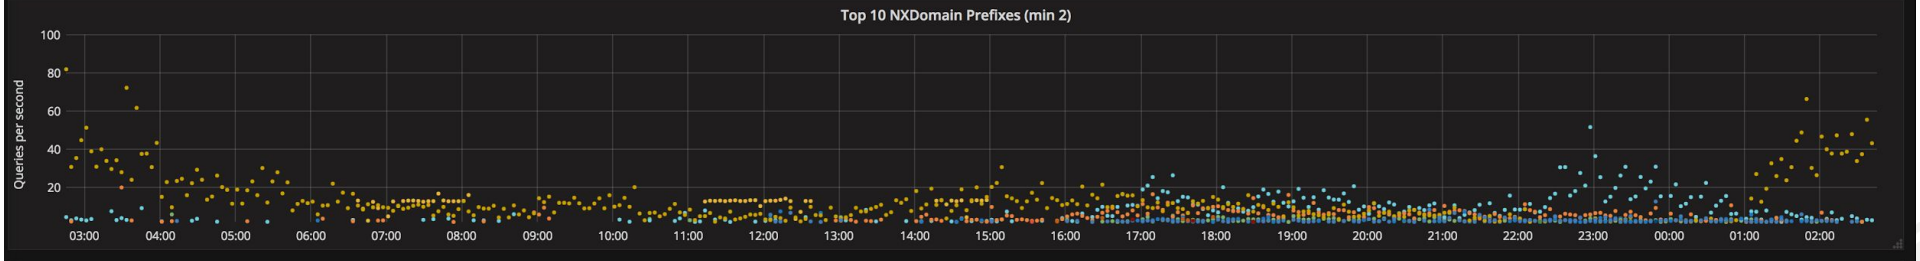

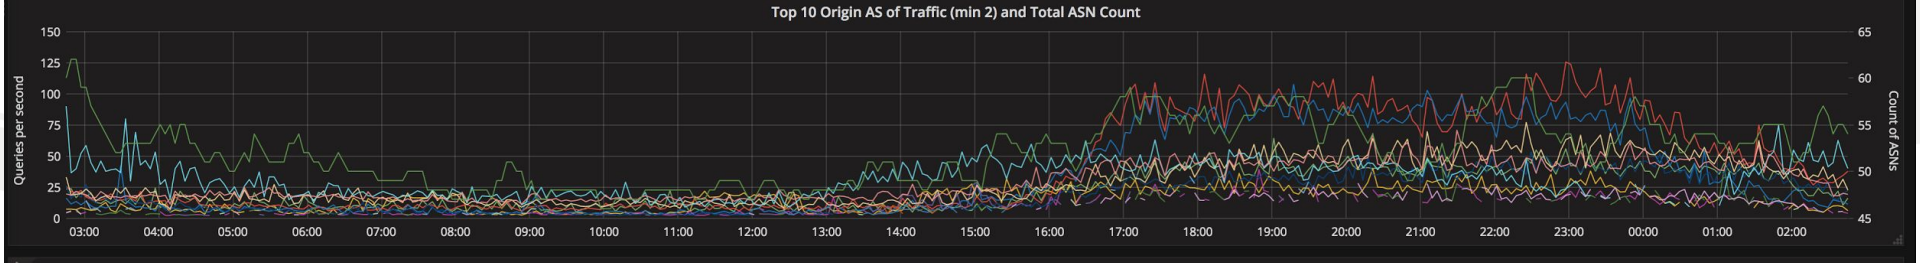

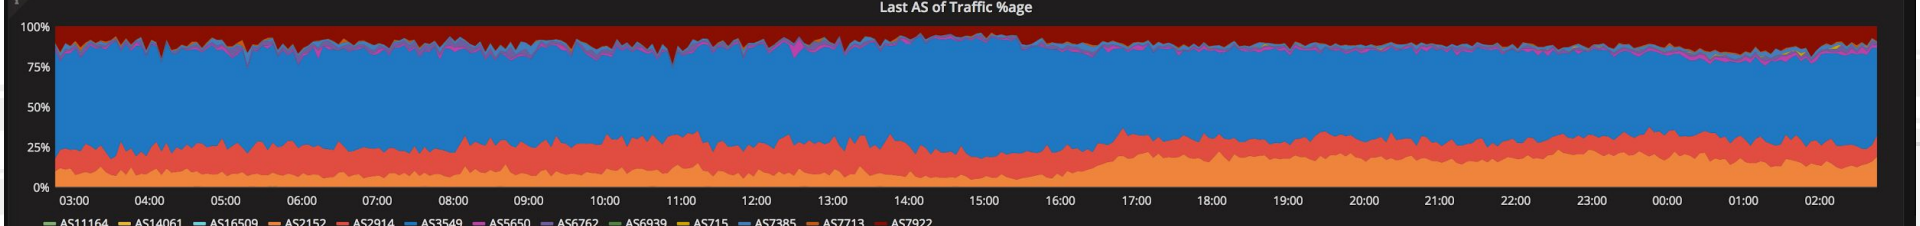

Results and current progress

DNSUni

- **Aggregation at edges**
  - We need to look at the whole feed of DNS answers
  - Memory caching at edges
  - Disk caching at central pop
- **Still too much data**
  - Need for graceful restart process at edges
  - Which means discarding some uniques queries
- **Feed of unique DNS records**

Conclusion/Lessons Learned

- **Specific architecture are a benefit in specific conditions**
  - Allows edge distillation/aggregation of data
  - Lose some ad-hoc ability; gain speed
  - Gain in scalability and decrease in cost
- **TSDBs**
  - Can also be used for complex monitoring needs
  - Names are important; decide between metric names and tags on the first version.
- **Design the codebase so it's flexible enough**
  - Reusable components
  - Not too much

Future projects

- **Future work for TSDB ingestion in the same model:**
  - Alternate ingestion protocols (DNSTAP)
  - Domain name monitoring (volume stats for zones)
  - Selected PCAP data into TSDB that is not specific to DNS using same summarization model

- **Future work for metrics management/display:**
  - The neverending dashboard creation process
  - Alerting/integration with monitoring apps that have “AI”

Thank you

Alexis Fasquel
Software Engineer
Packet Clearing House
alexis@pch.net