dnstap
(brief intro and update)

Merike Kaeo
merike@internetidentity.com

dnstap – What is it?

- High speed DNS logging without packet capture
  - Encoding uses Protocol Buffers
    - Binary clean
    - Efficient encoding
    - Extendable
    - Implementations available for many programming languages
- Schema file:
  - https://github.com/dnstap/dnstab.pb/blob/master/dnstap.proto

Why create dnstap?

- Frustration with certain limitations in packet capture approach to passive DNS replication
- Issues found in DNS-intensive analysis of pcap data
  - Bailiwick reconstruction
  - UDP fragment reassembly
  - UDP checksums
  - TCP stream reassembly
  - DNS query/response matching

Two specific use cases

- Query Logging
  - Make it faster by eliminating bottlenecks like text formatting and synchronous I/O
- Passive DNS replication
  - Avoid complicated state reconstruction issues by capturing messages instead of packets
- Able to support both use cases with the same generic mechanism

dnstap – Some of the how

- Add a lightweight message duplication facility directly into the DNS server
  - Verbatim wire-format DNS messages with context
- Use a fast logging implementation that doesn't degrade performance
  - Circular queues
  - Asynchronous, buffered I/O
  - Prefer to **drop** log payloads instead of **blocking** the server under load

dnstap – Message Types

- Present
  - Stub {Q,R}
  - Authoritative {Q,R}
  - Resolver {Q,R}
  - Client {Q,R}
  - Forwarder {Q,R}
- Prospective
  - RRL bucket {Start,End}
  - Zone Transfer in {S,E}
  - Zone Transfer out {S,E}
  - Cache Purge (LRU)
  - Cache Expiry (TTL)

dnstap-enabled DNS server

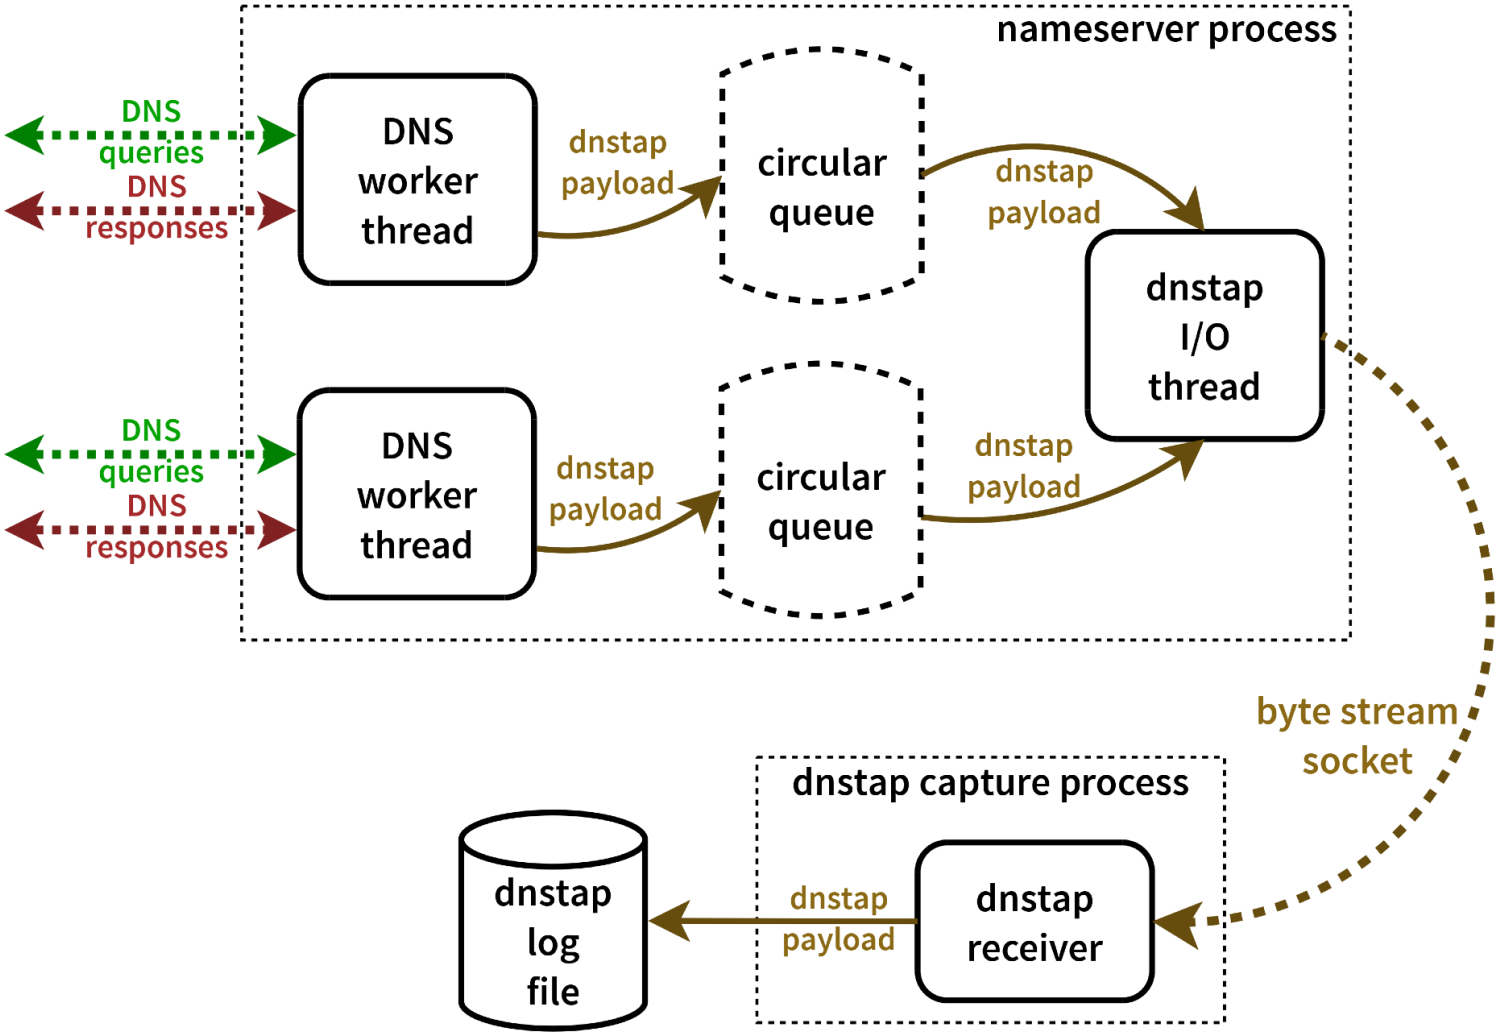

dnstap Components

- Flexible, structured log format for DNS software
- Helper libraries for adding support to DNS software
- Patch sets that integrate dnstap support into existing DNS software
- Capture tools for receiving dnstap messages from dnstap-enabled software

Advantages of dnstap

- Extensible and can offer mechanism into more detailed nameserver instrumentation
  - A way to easily audit the records that have been received **and accepted** by a recursive DNS server into its cache
  - Visibility into a new NS RRset overwriting an already existing NS RRset for the same name
- Vendor Neutral
  - Implementing dnstap support in multiple pieces of DNS software at once to make sure the same dnstap message means the same thing regardless of which software generated it.

Continued Outreach

- Speaking with varying DNS vendors to gauge receptiveness to integrate dnstap support into their software
- Initial proof of concept with Unbound
- CZ.NIC has already implemented dnstap support in the knot development tree
  - https://gitlab.labs.nic.cz/labs/knot

Feedback Wanted

- DNS operators interested in running their DNS servers with dnstap features enabled
- DNS implementors interested in adding dnstap features to their implementations
- DNS researchers interested in consuming dnstap-formatted data
- All feedback will be used to help set development priorities…..

Additional Information

- http://dnstap.info [definitive source]
  - dnstap.info/slides/dnstap_nanog60.pdf
  - dnstap.info/slides/dnstap_oarc2014_warsaw.pdf
- Robert Edmonds (edmonds@fsi.io)