Passive DNS Collection and Analysis The 'dnstap' Approach

Dr. Paul Vixie, CEO
Farsight Security, Inc.
2014-01-16 – Charleston, SC

Importance of Measuring DNS

- High volume low latency datagram protocol
  - `sie-xyzzy1  547,128,709,757 bytes,  80 sources ( 82%)`
  - `sie-xyzzy2   40,787,371,148 bytes, 141 sources (  6%)`
  - `sie-xyzzy3   21,650,049,219 bytes,  12 sources (  3%)`
- Enables almost all other network flows
  - A, AAAA, MX, NS, SRV records
- Traffic analysis: NetFlow vs. DNS
  - NetFlow tells you “what”
  - DNS tells you “why”

Challenges of Measuring DNS

- Historically, turning on logging in a DNS server slows it down to the speed of the file system
  - Operationally, measurement loss is always better
- So, success in DNS measurement has come from an asynchronous approach – BPF/pcap
  - NCAP (2006) – look for authoritative responses, reassembling UDP datagrams as necessary (EDNS)
  - NMSG (2009) – like NCAP but wants to also see requests, and log complete DNS transactions

Passive DNS Data Flow

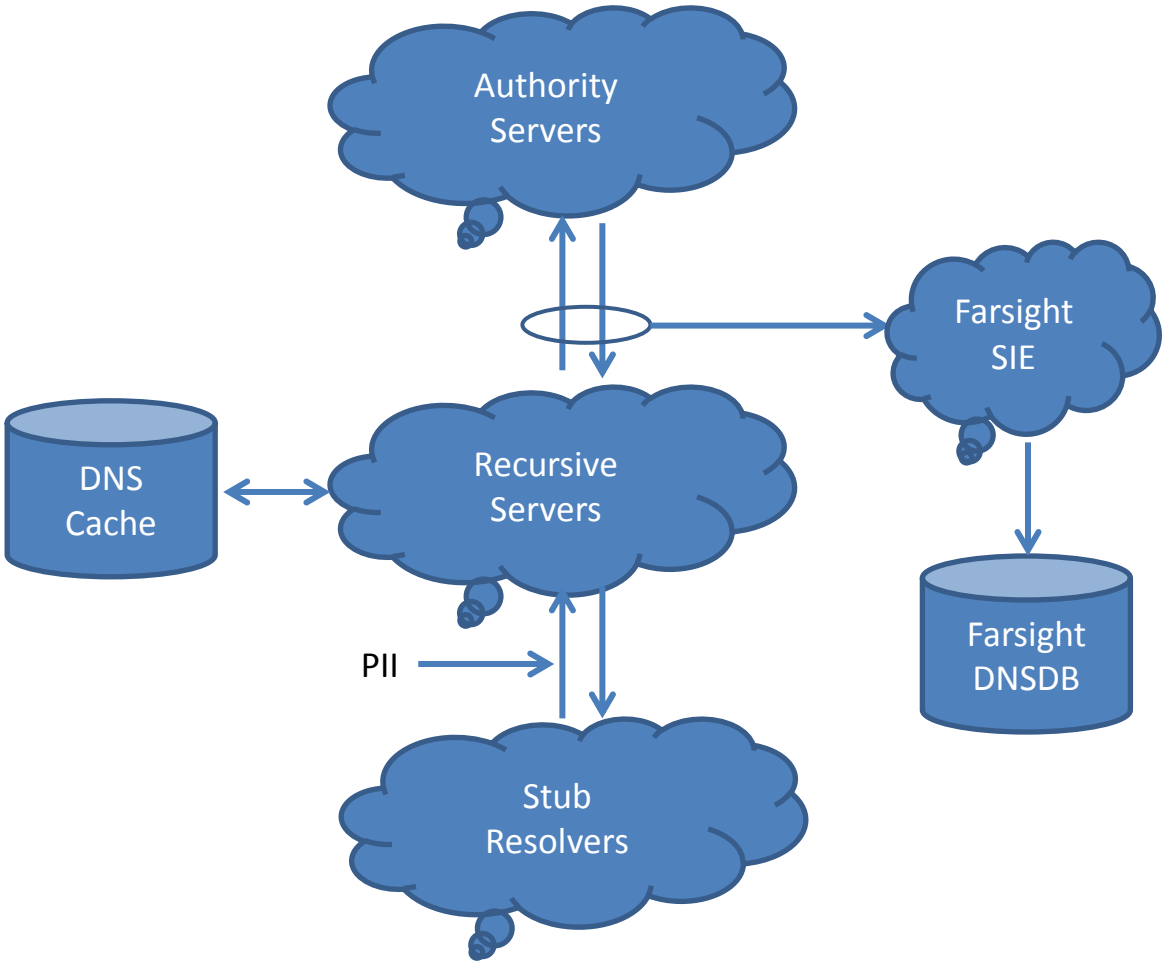

Problems with NMSG Approach

- Blind to off-the-wire events like cache expiry due to DNS TTL, cache purge due to LRU.
- Meaning is not tagged – NMSG receiver has to impute stub vs. cache miss transaction type.
- Currently blind to TCP/53 – noting that there can be many transactions per TCP/53 session.

Enter ‘dnstap’ (DNS Tap)

- Server-embedded
- TCP output streams
- Reliable front-loss
- Transactions, events: all tagged
- Apache licensed

‘dnstap’ Architecture

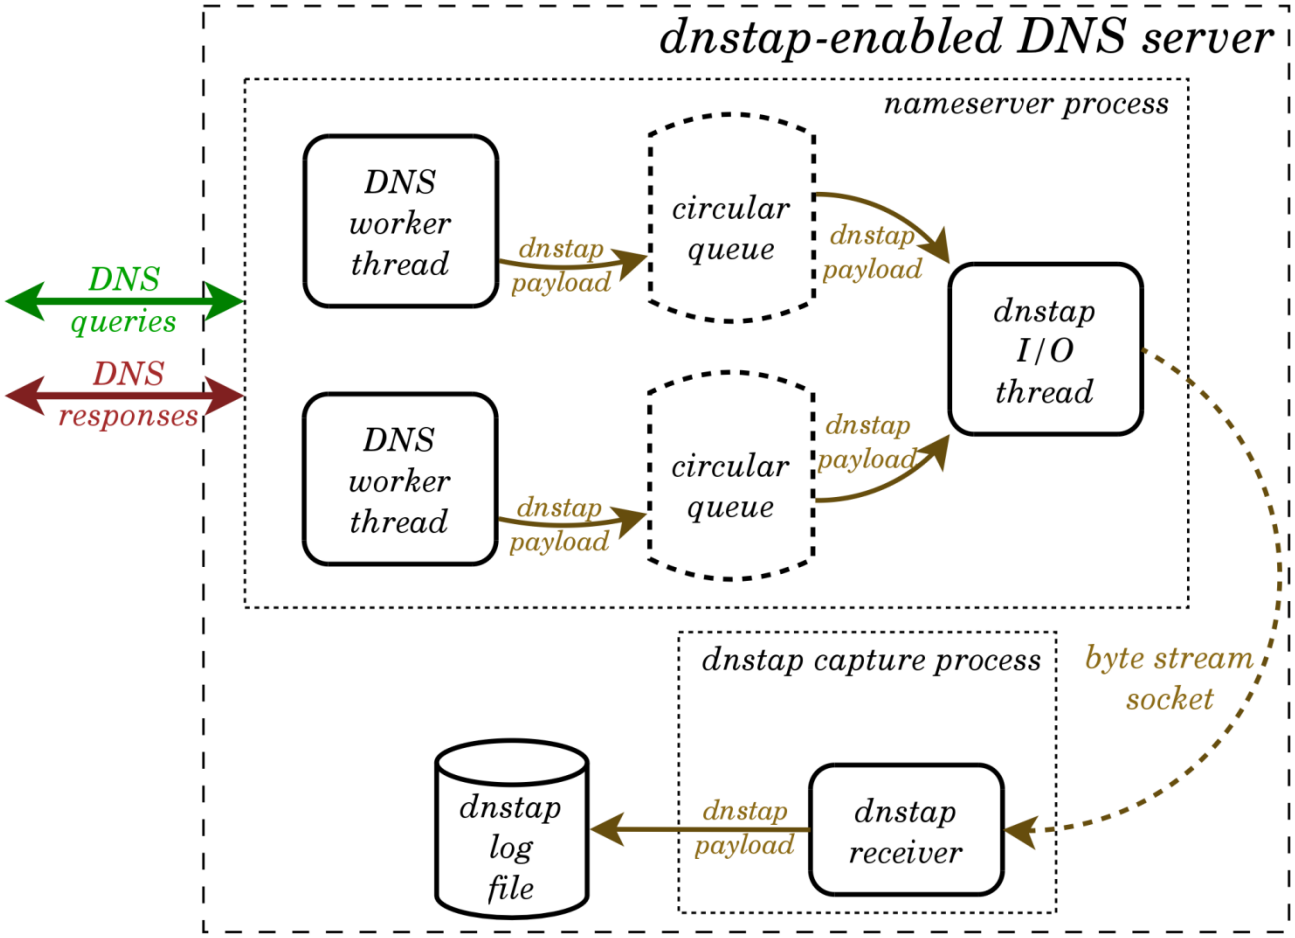

‘dnstap’ – Server-Embedded

- ‘dnstap’ messages are generated from within DNS implementations, via instrumentation
- So, no UDP fragment reassembly, no matching of on-wire queries with on-wire responses, and no worries about TCP/53
- We have this working in ‘unbound’ today
- ‘nsd’, ‘knot’, ‘powerdns’ and BIND: coming

‘dnstap’ – TCP Output Streams

- A ‘dnstap’ stream is a reliable byte stream
- So it can be a file, or a TCP session
- (Files? Some people really do like ‘rsync’)
- TCP means we won’t use >80% of channel
- TCP is easier on (inevitably) stateful firewalls
- Yet, TCP is unfortunately *very* (too) reliable

‘dnstap’ – Reliable Front-Loss

- TCP protocol vs. “Sockets API”
  - Nonblocking UDP socket rejects full datagrams
  - Nonblocking TCP socket rejects overflow octets
    - Which breaks “framing” unless sender keeps state
    - But we want total message loss in this case!
    - And we want such messages dropped *early*
- Solution: ‘dnstap’ writer thread
  - Lockless SP/SC ring buffer
  - ‘dnstap’ socket is blocking, so, thread can block
  - Reliable front-loss occurs when ring buffer is full

Congestion (Thanks: Van Jacobson)

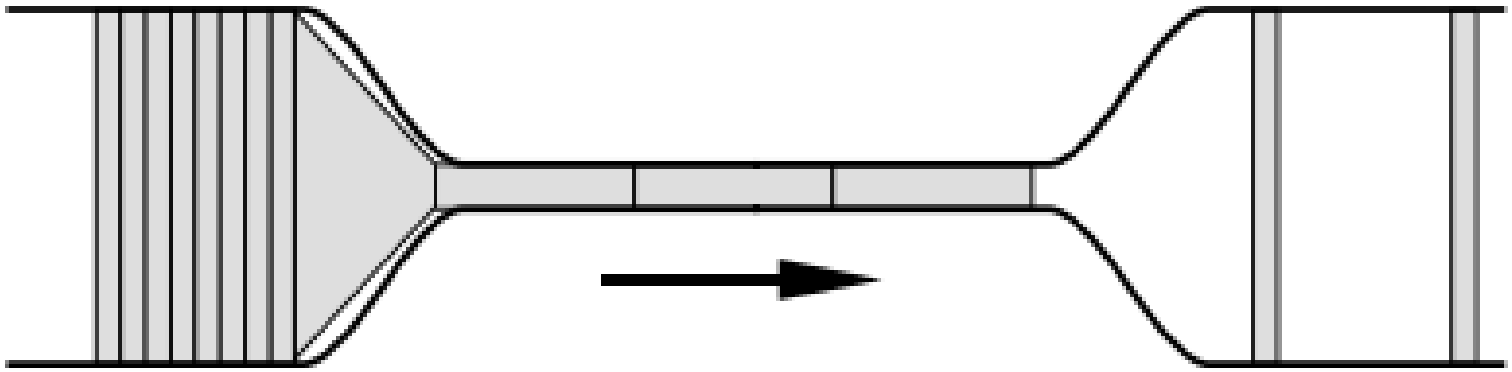

‘dnstap’ – Message Types

- Present:
  - Stub {Query, Response}
  - 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’ – Licensing/Packaging

- Using Apache Open Source License V2.0
- We love BSD/ISC license; AOSL2 is “better”
- Protocol, reference API, reference toolset
- Our commercial interest is: wide adoption
  - So, it’s all on GitHub (see http://dnstap.info/)
- We intend to patch all F/L/OSS DNS servers
  - ‘dnstap’ is structured as a copy-in, not a dependency, noting that it depends on protobuf-c

Context of DNS Measurements

- Farsight (was ISC) SIE – Security Info. Exchange
  - Commoditize security-relevant Internet telemetry
  - Channels for Passive DNS (raw, dedup'd, validated, filtered, chaff)
- Filtered output goes into DNSDB
  - Hierarchical MTBL (Google Sorted String Table)
  - RESTful API, JSON output
  - Stored everything from SIE since June 2010
- SIE and DNSDB are cash-free for nonprofit research/academia (pay us in data of like kind)

Passive DNS, SIE, DNSDB – Context

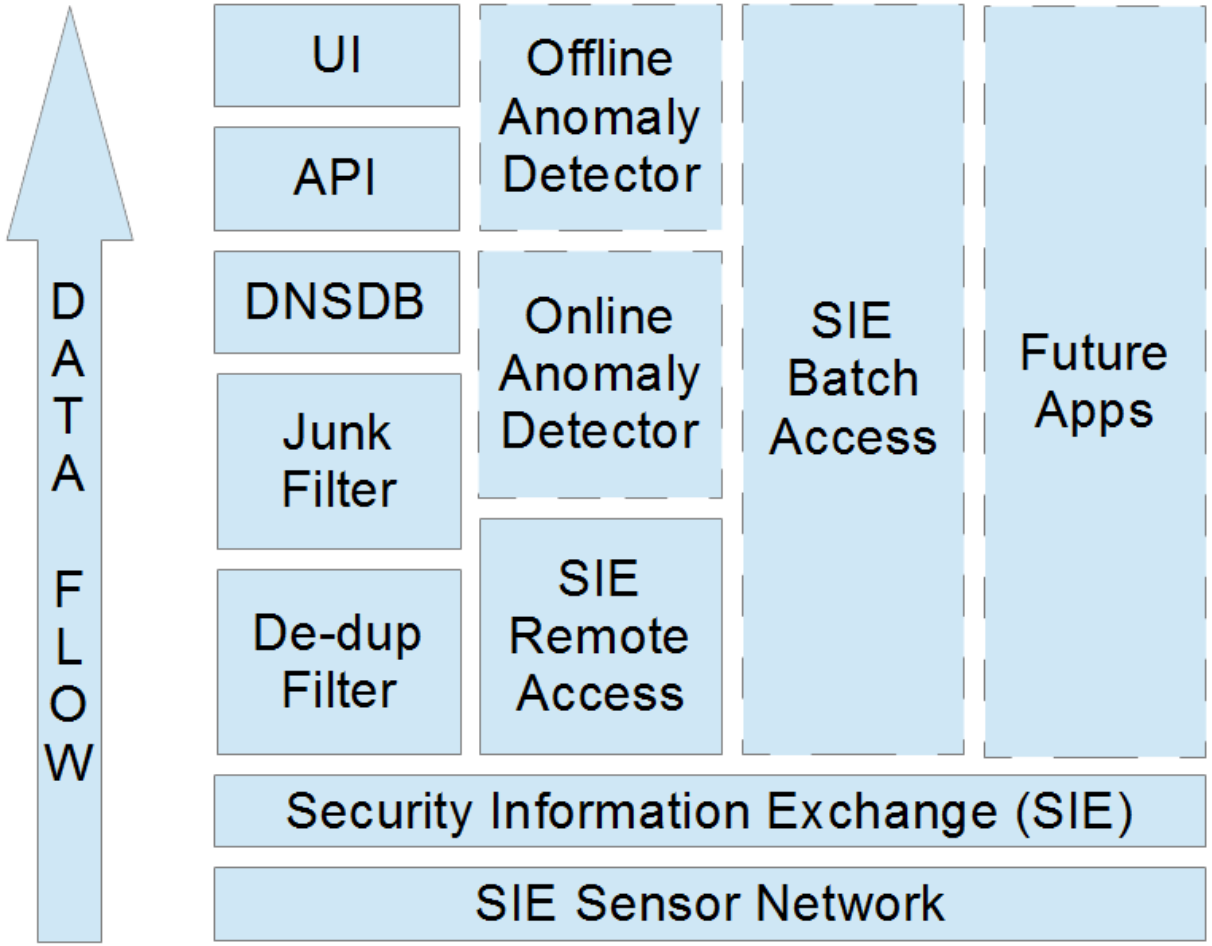

Demonstration

- SIE – nmsgtool, tcpdump
- DNSDB API – online “dnsdb_query” tool
- SRA – SIE Remote Access (pre-release)
- DNSDB UI – web user interface for LEA

Summary

- Passive DNS monitoring (NCAP, NMSG)
- ‘dnstap’ (coming during 2014)
- Worked example: SIE and DNSDB
- More Information:
  - http://dnstap.info/
  - https://dnsdb.info/
  - https://api.dnsdb.info/
  - http://github.com/farsightsec