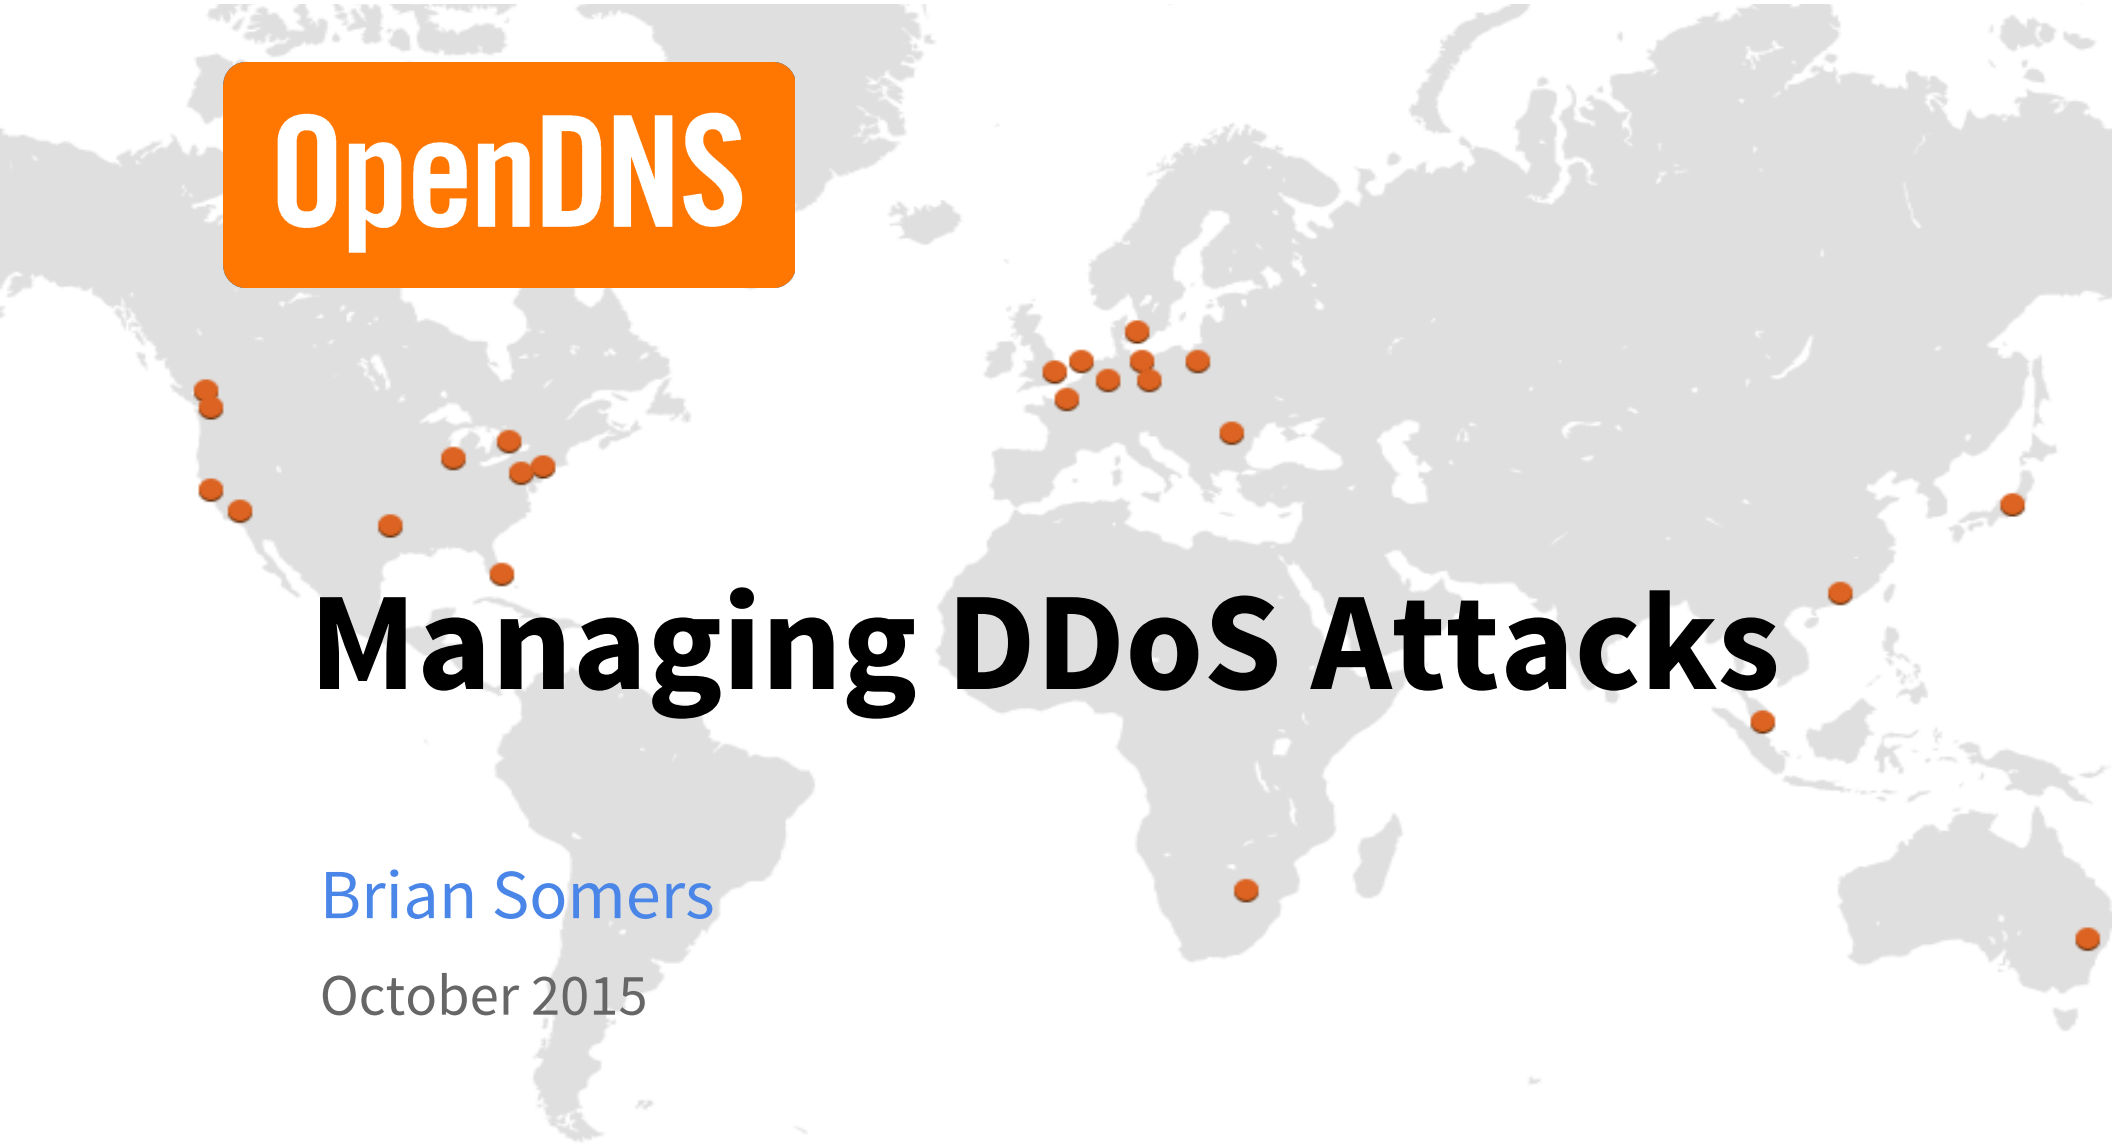

Managing DDoS Attacks

Brian Somers

October 2015

who am i

bsomers@OpenDNS.com

brian@FreeBSD.org

brian@Awfulhak.org

Brian Somers
OpenDNS Site Reliability Engineer

Responsible for the **Manageability**, **Scalability** and **Reliability** of the OpenDNS resolver infrastructure

DNS DoS

- Accidental Attack
  - Queries from one source repeated too often
- Amplification Attack
  - Falsified source (the target)
  - Biggest possible answer payload
- NXDomain Attack
  - Targeting a specific authority
  - Answer not likely to be cached
  - Uses a botnet **and** falsified source addresses

Rate Limiting

- Most basic limitation, handles Accidental and Amplification attacks
- Based on the source address
- Limiting becomes more aggressive if the source is persistent
- Open to abuse by faking the target's IP as the source IP
- Different limits based on client categorization, query type, domain categorization, response size, client customer status

Global Attack Identification

Our big data systems consolidate query statistics and identify possible attacks based on the following criteria

- At least 500 unique domains queried in a 10 second period

  (this is unusual except for public suffices and RBLs)

- Negative response ratio must be (95%) of which at least 30% must be **SERVFAIL**
- The longest domain is chosen

The domain-droplist

- Queries to domains in the domain-droplist are dropped; no response is sent to the client
- The domain-droplist is populated based on domains satisfying the Global Attack Identification criteria, and having an average of less than 100 queries per hour over the last two weeks
- The domain-droplist update is pushed globally in just a few seconds

The domain-freezelist

- Queries to domains in the domain-freezelist are served if a cache entry (expired or not) can be found
- Queries to domains in the domain-freezelist are served if the client can be identified (via EDNS0 or source IP)
- Queries to domains in the domain-freezelist are served over TCP
- The domain-freezelist is populated based on domains satisfying the Global Attack Identification criteria, and having an average of between 100 and 500,000 queries per hour over the last two weeks
- The domain-freezelist update is pushed globally in just a few seconds

Dropped Data

A single resolver can drop thousands of queries per second. Drops (magenta) peak at 4,780/second here.

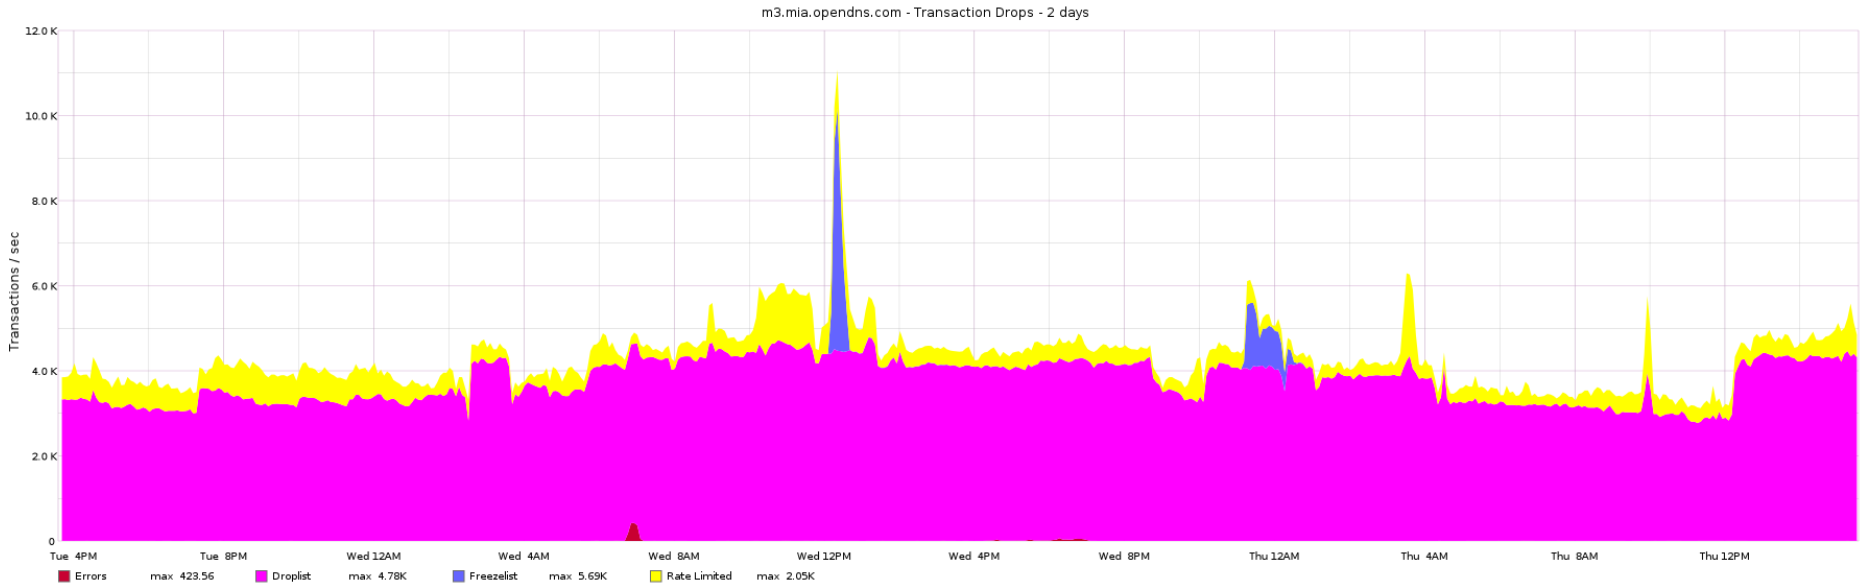

Transaction Ratios

As a ratio against “normal” traffic (purple), up to 40% might be dropped (green). Daily this is about 15 billion of 95 billion.

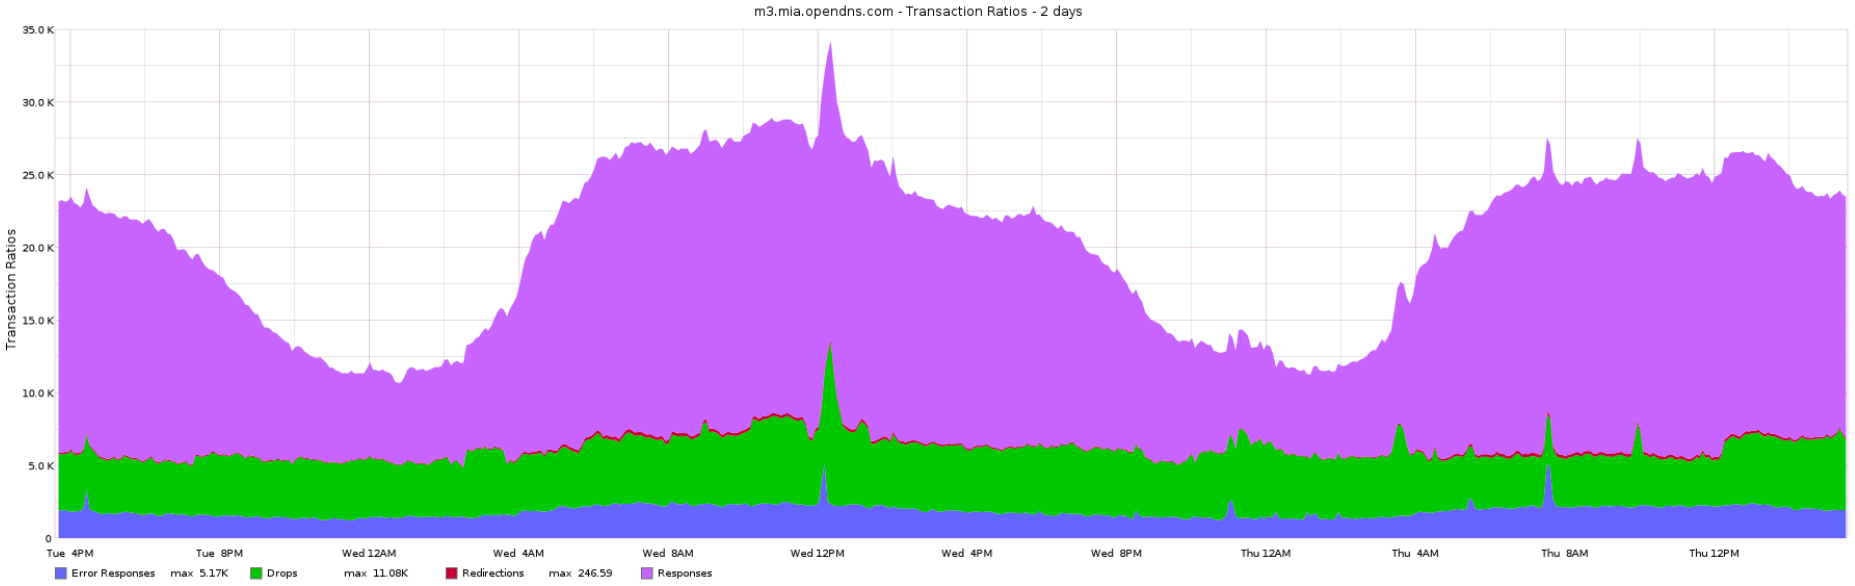

Authoritative RTT

We also avoid performing our own DoS against authorities

- Maintain RTT per control/qtype/auth-IP combination
- Candidate nameservers are used lowest RTT first
- Avoids querying latent authorities
  - May be under attack
  - May be "far away"
  - May be under maintenance
- If latency hits a predefined threshold, the authority is considered "down"
  - Put in timeout (40 seconds)
  - May cause immediate SERVFAIL responses to clients

NXDOMAIN Attack Footprint

DROPed (magenta) and

ratelimited (yellow) traffic

SERVFAIL (red) vs

SMARTCACHE (green)

SERVFAIL (salmon) vs

other non-zero rcodes

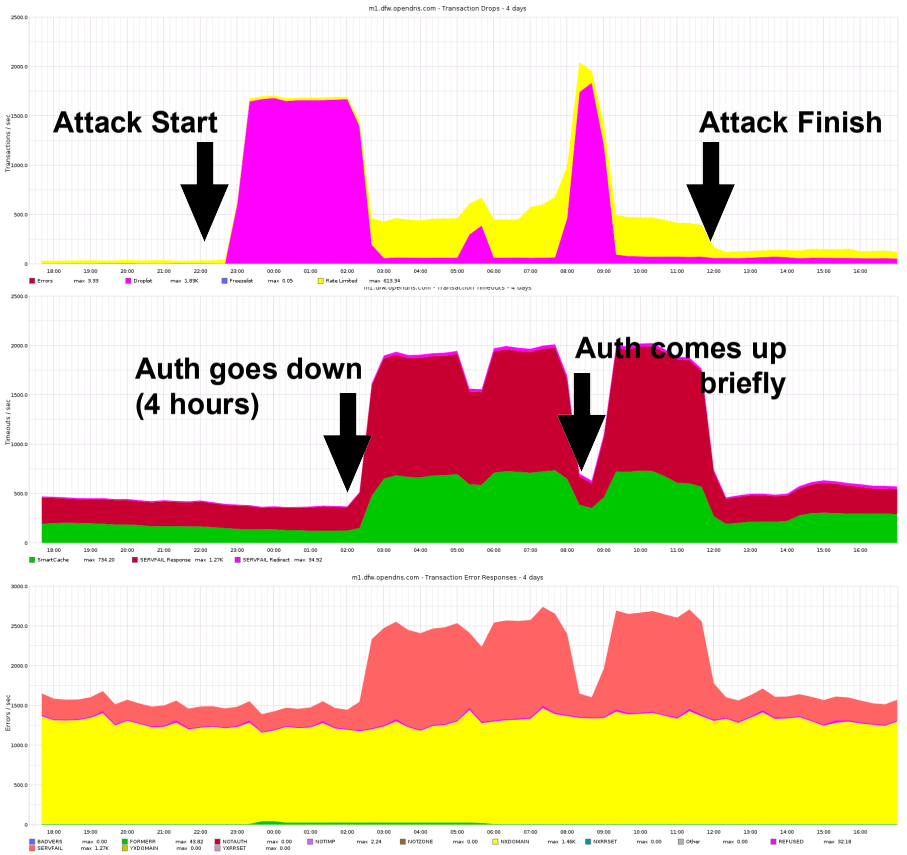

NXDOMAIN Attack Footprint

The attack is all additional traffic (green)

Getting SERVFAIL (brown) and Oversize (yellow) limiting

Resolver “suffers” keeping additional query context

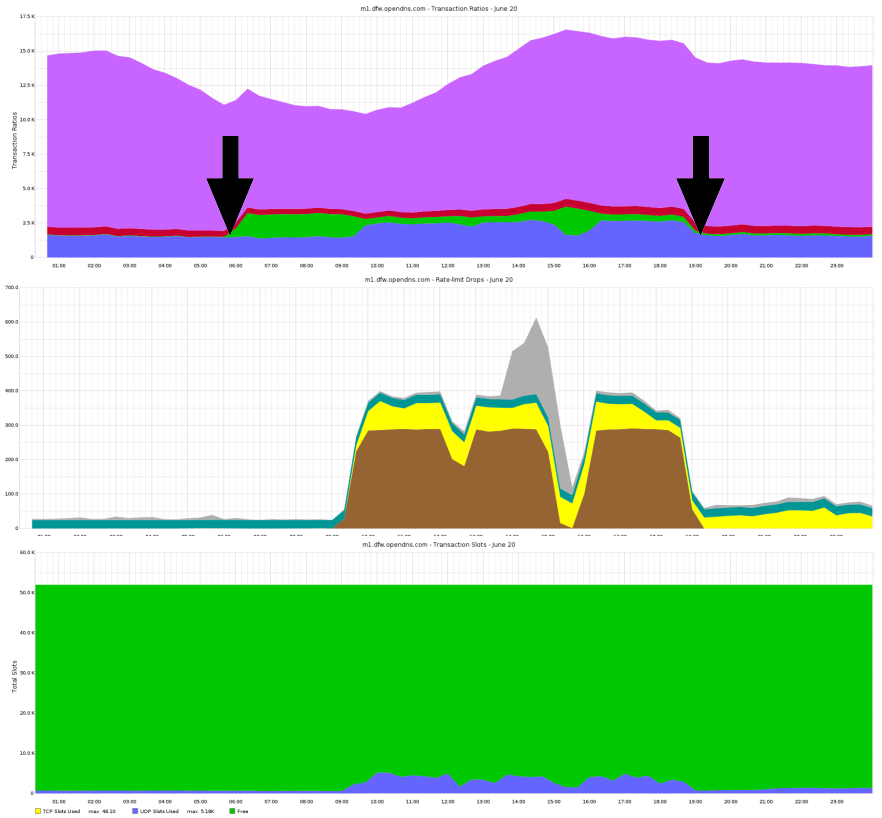

Future Technologies

Improved Freezelist “Thawing”

Improvements are in magenta

- UDP queries matching the freezelist
  - If in the main-cache (maybe expired), proceed as normal
  - If in the freeze-cache (small cache), send a TC
  - Store in the freeze-cache
  - Drop
- TCP queries matching the freezelist
  - Respond as normal
  - Close the connection after response is sent
  - Broadcast the answer to all resolvers in the datacenter

NXDOMAIN count

Store an NXDOMAIN count at the zone cut

- Z.mydomain.com stored as level1-nxdomain-count
- Y.Z.mydomain.com stored as level2-nxdomain-count
- X.Y.Z.mydomain.com stored as level3+-nxdomain-count
- **.X.Y.Z.mydomain.com stored as level3+-nxdomain-count
- Ratelimit based on per-zone *-nxdomain-count

Whitelist Labels

Labels such as {www,mail,ns}0?[0-9]? should be whitelisted

- Apply the whitelist to domain-freezelist at top level only
  - www01.target-domain.com is whitelisted
  - www01.ac84lsdlies.target-domain.com is not whitelisted
- Apply the whitelist to NXDOMAIN counts at level <= 3
  - mail.target-domain.com is not counted
  - mail.a.b.c.target-domain.com is counted