DNS Anycast Statistic Collection

RIPE 61 Measurement Analysis and Tools Working Group

18 Nov 2010

Edward Lewis

Neustar

What's so hard about reporting?

ed.lewis@neustar.biz

Collecting DNS Statistics (Generic)

- **The Technical Problem to Solve**
  - » Multiple remote sites with small limited local storage to store statistical data
    - – Lots of data to observe
  - » Central analysis point
- **Choices**
  - » Sampling, summing or packet capture
  - » Store & send or pre-process remotely
- **What is to be Learned**
  - » Activity by wall-clock and/or event
  - » Rough approximation
  - » Trends

ed.lewis@neustar.biz

It's not a few drips, it is a fire hose

- **Our servers see 12-20 billion queries per day**
  - » Rough estimates:
    - – Queries 80 bytes
    - – Responses 300 bytes
  - » That's pre-DNSSEC averages
- **Total size per day, tracking everything**
  - » Queries: 1.5 TB
  - » Responses: 5.5 TB
  - » Total traffic *compresses* to maybe 3 TB/day
  - » DNSSEC records don't compress as well

Sampling just 1% of data, still equals 30 GB/day (compressed)

ed.lewis@neustar.biz

Have to stick to what's important

Focus: cut down on traffic, to manage the analysis

- **QNAME and QTYPE**
- **Originating IP**
- **Minute granularity**

- **Which server answered and how well (fast)**

ed.lewis@neustar.biz

Accuracy

- An important factor
  - Is the collection accurate?
  - Is the analysis accurate?
  - Important for Monitoring, Analysis and Forensics
- More importantly - Does it agree with billing data?
  - Was the billing data accurate?
  - Engineer's "white-knuckle" question
- The test: Would you use this for billing too?

ed.lewis@neustar.biz

Neustar's Approach

ed.lewis@neustar.biz

High-level Design

- Each site has Capture Devices
- Data Routed to central storage devices - Netezza
- Oracle Business Intelligence Engine (OBIEE) performs reports against the Netezza device.

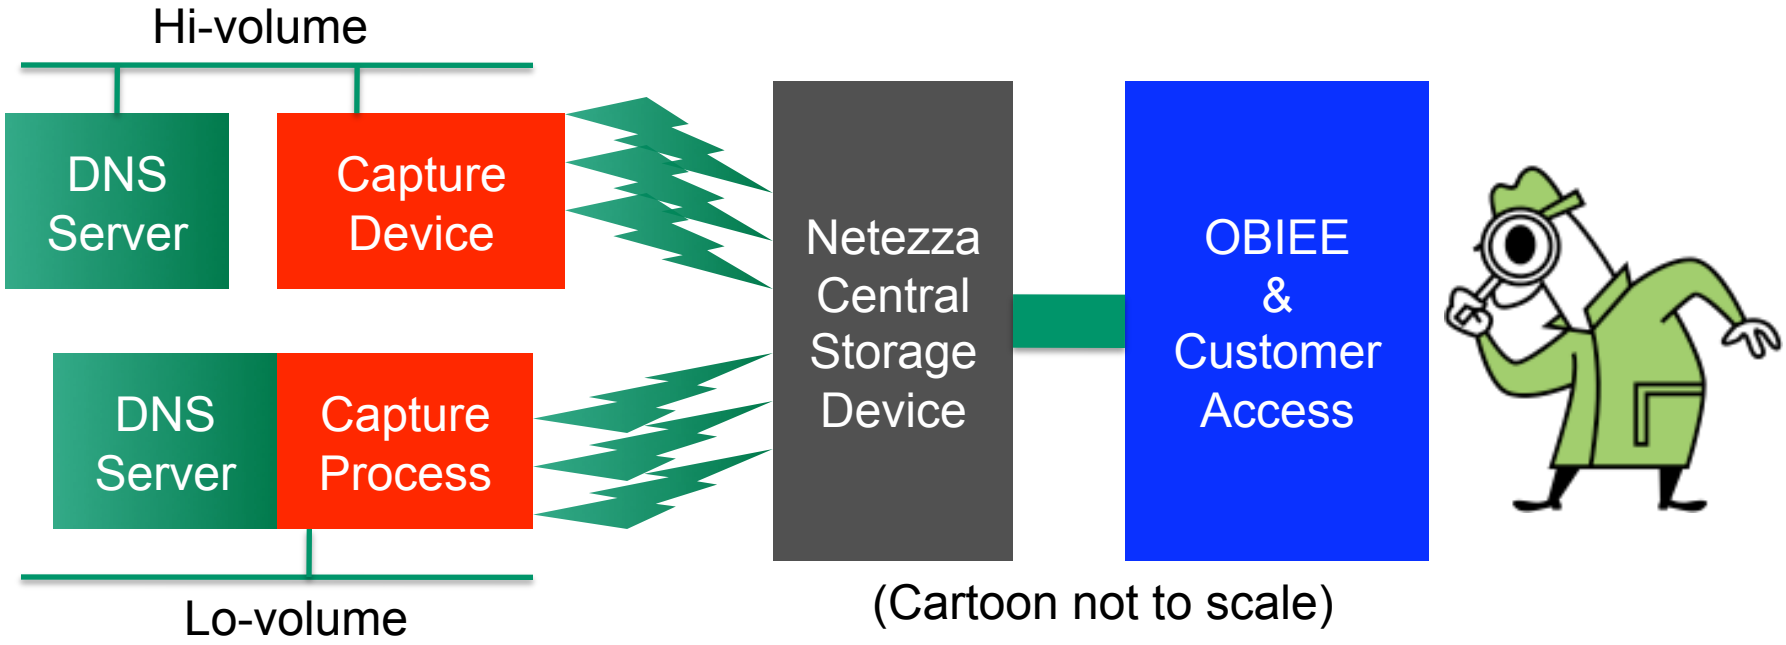

ed.lewis@neustar.biz

Capture, reporting, analysis approach

- **Collect and sum locally**
  - » High volume nodes use two in-line network taps
  - » Low traffic nodes use software on the DNS host
  - » Each query saved until matched to a response
  - » Aggregated data counted and compressed locally
- **Send to central repository every *second***
- **Central Netezza Database for analysis**
  - » Unique zones and QNAMES are loaded to dimension tables (about 1 billion unique QNAMES / month).
  - » Metrics loaded to fact table (~2billion rows/day).
  - » Map Source IP Address to Geographic Location

ed.lewis@neustar.biz

What data is captured

- **Per Minute, per node, query detail**
  - » Network Protocol *(IPv6 or IPv4, UDP or TCP)*
  - » Query Name and Type
  - » Response Code *(e.g., SERVFAIL)*
  - » Response Time Bucket *(<1, 10, 100ms, etc)*
  - » Complete NXDOMAIN traffic *(QNAME)*
  - » Unanswered Queries *(malformed, lame, etc.)*

- **Sampled Data**
  - » Track minimum of 1% of all traffic
  - » Query's source IP, full response message
  - » Will likely add logic to store sample for "irregular" queries

ed.lewis@neustar.biz

Compression is King

- **Netezza**
  - » Provides 6x compression of numeric data
  - » QNAMEs are saved in a separate table to maximize this
- **Messages as bitmaps**
  - » Encode a DNS message in a 64 bit integer
  - » UDP/TCP (1 bit), Yes/No: A, AAAA, NS, NXDOMAIN, ...
  - » Count how many times this appears in a minute
- **Minutes expressed as bitmaps**
  - » Many minutes are the same too
- **Relying on observed network behavior**
  - » Data can be compressed 30x to 40x compared to raw packet stream

ed.lewis@neustar.biz

Reporting Capability Detail

Reporting Capabilities			
	Pre -2010 Reporting	Standard Reporting	Advanced Reporting
Objects Under Management	Partial	X	X
Reporting Interval – Monthly	X	X	X
Reporting Interval – Daily		X	X
Reporting Interval – Hourly		X	X
Reporting Interval – Minute			X
Near Real Time Reporting			X
Total # of Queries	A, Z	A, Z, Q	A, Z, Q
TCP vs. UDP Queries		A, Z, Q	A, Z, Q
IPv4 vs IPv6 Queries by Zone		A, Z, Q	A, Z, Q
Queries by Query Type		A, Z, Q	A, Z, Q
Queries by Node Region	A, Z		
Queries by Source Country		A, Z, Q*	A, Z, Q*
Avg/StdDev Query Response Time		A, Z, Q	A, Z, Q
Queries by US State / CA Province / Zip Code			A, Z, Q*
Queries by Source IP			Q*
Queries by RCODE (Includes Errors)			A, Z, Q
Forecast Monthly Query Amount		X	X
Monthly Trending Reports		X	X
Download Reports	X	X	X
Easily View Trends		X	X

*Reports are based on Sample data

A = Account , Z=Zone, Q = QNAME

ed.lewis@neustar.biz

“Stories from the Road”

OR “Potholes we managed to hit….”

ed.lewis@neustar.biz

Queries and Responses

- **Is every query answered?**
  - » Can't just track the responses, must track queries and match to the responses.
- **How well (fast) are servers operating?**
  - » The DNS Server is only a portion of the site architecture. Must attempt to determine the time a query enters the site and determine the time the response leaves the site
- **Can the DNS Server do the measurement?**
  - » No way. It's what's being monitored
  - » *(joke) 100% of all dropped queries are never answered!*

ed.lewis@neustar.biz

Problems encountered in monitoring

- **QNAME population size**
  - » NXDOMAIN traffic for TLDs generates a lot of data, need to store it
  - » In an attempt to defeat cache poisoning, some recursive DNS service providers have appended random strings in front of queries to TLDs (inflating QNAME population)
  - » Certain Managed DNS service customers rely on random/undefined hostnames.
    - – *One large social networking site generates billions of unique QNAMEs.*

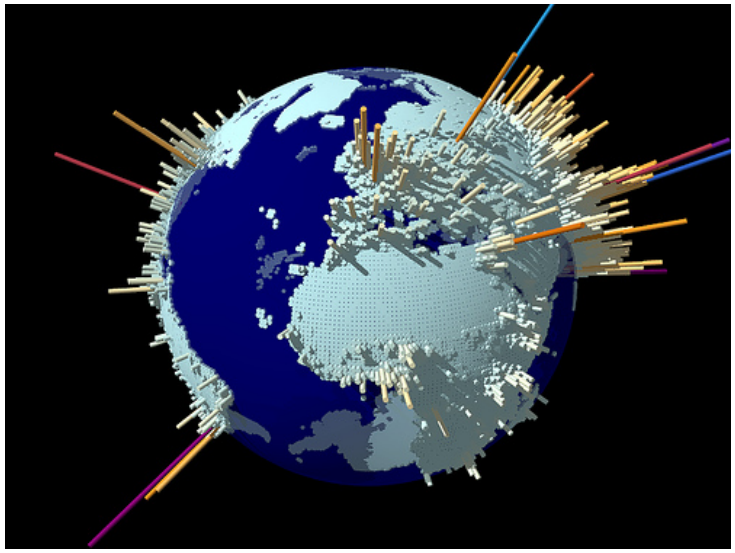

- **DDoS**
  - » Packet floods need to be mitigated, monitoring needs to record them but not fall to them

ed.lewis@neustar.biz

Things seen with this system

- **Sudden Traffic Growth Syndrome**
  - » Several causes – data gives chance to determine why
- **Rogue Recursive Servers**
  - » Some only "went rogue" on specific zones (big and small ISPs)
- **Routing issues**
  - » Matching answering anycast node with the source country
- **DDoS**
  - » Small DDoS attempts/attacks are noticed
- **The World isn't That Big**
  - » Top 1000 Source IP Addresses perform half of DNS queries

ed.lewis@neustar.biz

Some Stats

UDP	TCP	IPv4	IPv6
99.936%	0.014%	99.76%	0.24%

A	AAAA	CNAME	DNSKEY	MX	NS	PTR	TXT
75.01%	9.3%	0.10%	0.02%	10.34%	0.35%	3.2%	0.5%

Queries arriving via IPv6: 0.24%

Contrast to the nearly 10% of queries for AAAA

ed.lewis@neustar.biz

“There are no mistakes, only lessons. Growth is a process of trial and error.”

Where do we go from here?

ed.lewis@neustar.biz

2011 Roadmap

- Response by Server in Detail
- Alerts for Account Level Changes
- View/Download Raw Sampled Data
- Comparison Graphing
- Query Type Drill Down
- Additional Filtering
- Scheduled Emailing of Reports
- Access to Reporting Data via API
- Interactive Graphing
- Map view of Geo data
- User defined Ad hoc reporting

ed.lewis@neustar.biz

Questions?

ed.lewis@neustar.biz