DITL CRUNCHING FOR GTLD NAME COLLISION STUDY

Jim Reid
RTFM LLP/Interisle LLC

INTRODUCTION

- Why, When & How?
- Hardware choices
- Software choices & trade-offs
- Problems
- Results/Findings
- Possible next steps

BACKGROUND

- ICANN concerned about potential problems from new gTLDs clashing with existing *ad-hoc* use of these in domain names, "private" name spaces and certificates
  - Some anecdotal evidence, but no hard data
- Study approved by ICANN board in mid May 2013
  - Is there a problem?
  - If so, how big is it?
  - What risk mitigation frameworks could be applied?

TIMING

- **VERY** Ambitious!
- Gather & analyse data
  - First find out how best to do that and what resources can be brought to bear
- Write a report for Durban ICANN meeting in mid-July
  - => about 6 weeks away
- Got even scarier once the scope of the data crunching became apparent

‘NUFF SAID....

KICK-OFF

- Preliminary discussions took place at RIPE66 in Dublin
  - Many RSOs present, OARC meeting too
  - How best to get data and process them
- Use the DITL datasets at OARC
  - Only practical way to get access to suitable data
  - Simple, quick solution to privacy and data protection concerns
  - RTFM LLP became an OARC member :-)

OBJECTIVES FOR DNS ELEMENT OF THE STUDY

- Count how often new gTLDs appear in root server traffic
  - Are these requests localised or diffuse?
  - Proper resolving servers or from forwarders/stubs?
  - How does this compare to traffic for existing TLDs?
- How often do new gTLD labels appear elsewhere in QNAMEs?
  - Where do they appear?
- For bonus points, look at big resolver operators' traffic

INITIAL SCOPING

- Helpful advice and software from Netnod
- Got access to elderly box, **`an1.dns-oarc.net`**
  - 2-core 1Ghz Opteron, 2GB RAM, limited local disk
- Did some prototyping with **`packetq`**
- Some nasty shocks:
  - ~1000 new gTLDs found in a sample of the DITL pcaps
  - 1 pass over the 6TB of DITL pcaps for 2012 would take at least 2 weeks on this system: far too long

CAIDA TO THE RESCUE

- Lot of uncertainty over what other hardware could be provided:
  - Could anything be ordered, delivered and set up in time?
  - Maybe NFS mount the datasets into the cloud somewhere?
    - Throw a bazillion CPUs at the problem
- Found out CAIDA had a server which could be made available
  - 8-core 2GHz Xeon, 7TB of scratch disk space
  - Running 5-6yo version of FreeBSD
  - I pass over a year's DITL data would take less than a week

SOFTWARE CHOICES - 1

- Got a custom version of **`packetq`** from Netnod
  - SQL-like language for crunching through pcap files
  - Mostly counted things: QTYPEs, QNAMEs, source addresses
    - Special hooks to recognise existing and proposed TLDs
- Could drive all cores flat-out simultaneously
- Not so good for label position counting/checking though
  - 1 week of CPU time for each N-th level label to inspect

SOFTWARE CHOICES - 2

- Use **`tcpdump`** & **`fgrep`** for a second pass over the pcaps
  - Generated text files containing pretty-printed DNS requests where any label matched a proposed gTLD
    - "Only" several GB of text files to then analyse
- **`awk`**-based scripts chugged through these text files to do label position and source address prefix counts
  - Sometimes tripped over bad input data because of malformed (-ish) queries, e.g. ***`foo.bar.tld .`***

GENERAL APPROACH

- Split the ~250,000 pcap files for each year into 8 equal chunks
- Run script over each pcap as an "atomic" operation
  - Generate unique output files for each input file
    - Merge or aggregate these interim files later
  - Could process files by hand if bugs/corner cases pop up
  - No locking/synchronisation issues
  - Just keep crunching, never stop or go back
  - Flag errors as corner cases, but don't allow these to get in the way or complicate the scripting

TRIPLE DATA DISTILLATION

- 1: reduce terabytes of raw data to *O*(gigabytes) of rough results
- 2: distill rough results to *O*(megabytes) of refined results
- 3: feed refined results into spreadsheets and PHP-based tools for statistical analysis
  - Summary results analysed in more detail by Interisle
    - Some sampling done too
  - Interisle drew graphs and compiled tables for final report

WHY NO DATABASE?

- Couldn't realistically prototype/calibrate this in time
- Far too many unknowns
  - How big would the database(s) be?
    - What's the optimal size of the tables and indexes?
  - How long would it take to populate the database(s)?
    - Locking/synchronisation with 8 CPUs in parallel
  - How long would SQL queries take to run?
  - What if the database got corrupted or a scratch disk died?

WHY TCPDUMP?

- It'll be faster than **`perl`** or **`python`** or...
- Newest DNS tools need newer **`perl`**/**`python`**/***`whatever`*** versions than the CAIDA box had
  - Too many unknowns - install/software configuration hell
  - Needed certainty when results could be expected
- Just generate plain text output
  - Run **`awk`** scripts (mostly) over those files, 8 at a time
  - Merge/aggregate the results

HOW THE CRUNCHING GOT DONE

- 2 passes over both 2012 and 2013 DITL RSO pcaps
  - First pass counted TLDs using packetq
    - Took about 2 CPU-months (1 week of elapsed time)
  - Second pass got **`tcpdump`** to pretty-print packets where QNAME contained a new gTLD label
    - Took nearly 2 CPU-months
    - **`awk`** scripts took about 1 CPU-week to analyse label positions, count source address prefixes

TIMELINES

- Crunching the 2012 DITL pcaps took around two weeks of elapsed time
  - Finished in mid-June 2013
- Most of the 2013 DITL pcaps had arrived at OARC by then
  - Processing started on them as soon as the 2012 crunching was finished
  - Back-end "polishing" of results done by Interisle

AN UGLY GOTCHA

- Some **`packetq`** output from 2013 data was wrong:
  - Essentially null files were produced
- Duane Wessels explained some raw pcaps had used 802.1q VLAN tagging
  - Hadn't been tidied up at OARC at that point
- **`packetq`** treated 802.1q link-level header as payload
  - An off by 4 bytes error...
  - Henrik Levkowetz at Netnod fixed this very quickly :-)

FINDINGS

- Lots of power-law distributions
  - Small numbers of TLDs and source addresses (per TLD) accounted for most of the traffic
- **FAR** more traffic for proposed TLDs than gut feel suggested
  - Almost all new gTLDs were seen
  - Traffic for **`.home`** and **`.corp`** was particularly high
- Pretty much none of that DNS traffic was localised (enough)
- Some interesting/unexplained traffic patterns

FOR FURTHER ANALYSIS?

- Probable leakage from Active Directory and Bonjour
  - How will those end systems behave if/when NXDOMAIN becomes a referral response?
  - Some dynamic updates too....
- Lookups for MX and SRV records
  - Can't be coming from naive end users & applications
  - Something's been deliberately (mis)configured to look for these: what? why?
- Should be looked at in more detail

THE "SAFE" QUERY RATE THRESHOLD

- Lot of undue comment and attention on this
  - Ask ICANN, not Interisle
- The **`.bv`** and **`.sj`** TLDs are empty and unused
  - Nobody has a valid operational reason for querying them
  - Traffic volume they get seems a fair indication of the DNS background noise level as seen in root server traffic
- This is only one metric out of many and might well not be the most significant one for assessing new gTLD "safety"

LESSONS LEARNED

- The old tools from the 80s should not be overlooked
  - They can still do useful stuff
  - Fewer unknowns & dependencies, especially on an old OS
    - Generally better than the Swiss army penknife approach that **`perl`**/**`python`**/***`whatever`*** seem to typify
    - Avoid crufty bloatware and its rat-hole of dependencies
- Always remember Ken Thompson's advice: "when in doubt, use brute force"

ACKNOWLEDGEMENTS AND THANKS

- kc claffy and Daniel Anderson at CAIDA
  - Simply couldn't have done the work at that time without access to their hardware
  - Box died shortly after the crunching stopped...
- Henrik Levkowetz at Netnod
  - For tweaking and supporting **`packetq`**
  - Also did some sanity checking of early results
- OARC, especially William Sotomayor, for logistical support

QUESTIONS?