RIR delegation reports and address-by-economy measurements

DNS-OARC Workshop
25 July 2005

George Michaelson
APNIC
ggm@apnic.net

Summary

- All RIR produce daily reports on resource allocations and assignments
  - Consistent format, single source for ASN, IPv4 and IPv6
  - Simple (CSV friendly) format common to all RIR
  - Easy to combine to a global view
- Provides overview of resources by
  - Economy of registration
  - Size of delegation
  - Date of delegation
- Please use them!
  - With caveats to applicability and accuracy…
- Extensions coming..

‘delegated’ file format

- Produced daily
- Plaintext file, # comment, with external checksum and GPG signature files.
- Extensible by extra fields at end of line
  - Skip-fields to allow column-summing in spreadsheets etc
- Version tagged, with summary & range check data inline

```
2|apnic|20050725|13791|19850701|20050722|+1000
apnic|*|asn|*|2055|summary
apnic|*|ipv4|*|11345|summary
apnic|*|ipv6|*|391|summary
apnic|JP|asn|173|1|20020801|allocated
apnic|JP|ipv4|58.0.0.0|131072|20050106|allocated
apnic|JP|ipv6|2001:200::|35|19990813|allocated
```

Delegated file location(s)

`ftp://ftp.<rir>.net/pub/stats/<rir>/delegated-<rir>-latest`

ftp://ftp.apnic.net/pub/stats/apnic/delegated-apnic-latest

- We mirror each other, but its best to fetch from the source.
  - URL form is consistent at all RIR
- 'current' is the head file,
  - archive rolled to yyyy-mm-dd versions, compressed
- Also publish .md5 and .asc checksum/signatures in separate files.
- We make a fileset for IANA data
  - Completeness (pre-RIR direct allocations, reservations)
  - Shows downward delegation dates for Registry blocks
  - Ideally, would like IANA to publish themselves…
- Joint file production a candidate for NRO website

Caveats on Data

- Some dates are unreliable
  - Lack of data for sri-nic, ddn-nic, pre-RIR assignments.
  - Record updates can change apparent date
    - Splits, m & a, extension into reserve
- Some economy tags are unreliable
  - mis-marked (eg USAF Airbases overseas)
  - Transfers not completely documented
  - Increasing use of ‘aggregate’ codes EU/AP/ZZ
- IANA data is ‘our view’ of their data
  - Eg network 7.0.0.0/8 status is not well documented

Caveats on Applicability

- We believe size of resource is good
  - Please report errors!
- We believe date of delegation is mostly good
  - Some Interpolated dates.
    - pre-RIR assignments, handed out in almost-linear order
  - Some missing data (IANA mainly)
- Economy is economy of registration
  - Nets can be used anywhere worldwide
  - Some agencies (eg IBM GSA) have extremely large global footprint with historically assigned nets

Some example ways to use

1. Timelines (resource by date)
   - eg “The BGP movie”
     - Uses distinct generation versions of files
     - Semi-brute force approach
     - Combines BGP routing data
   - How much resource is active today? Measures
     - Opportunities for inter-generation comparisons
     - ‘rate of transaction’ reports
     - Trends analysis..
2. Where did my packet come from?
   - Reverse-DNS query by economy
   - Applicable to logfiles, tcpdump data
   - First-approximation measure (see caveats)

Timelines

- Using the files as the prime data itself
- Little or no data analysis required
- Approx 100k records in 6 sources per day, spanning 1986-2005
- Sort & map delegations into 2-D barchart, one per day as .JPG
- Animate with overlays
  - Netpbm, ploticus, perl, sed, awk, sweat

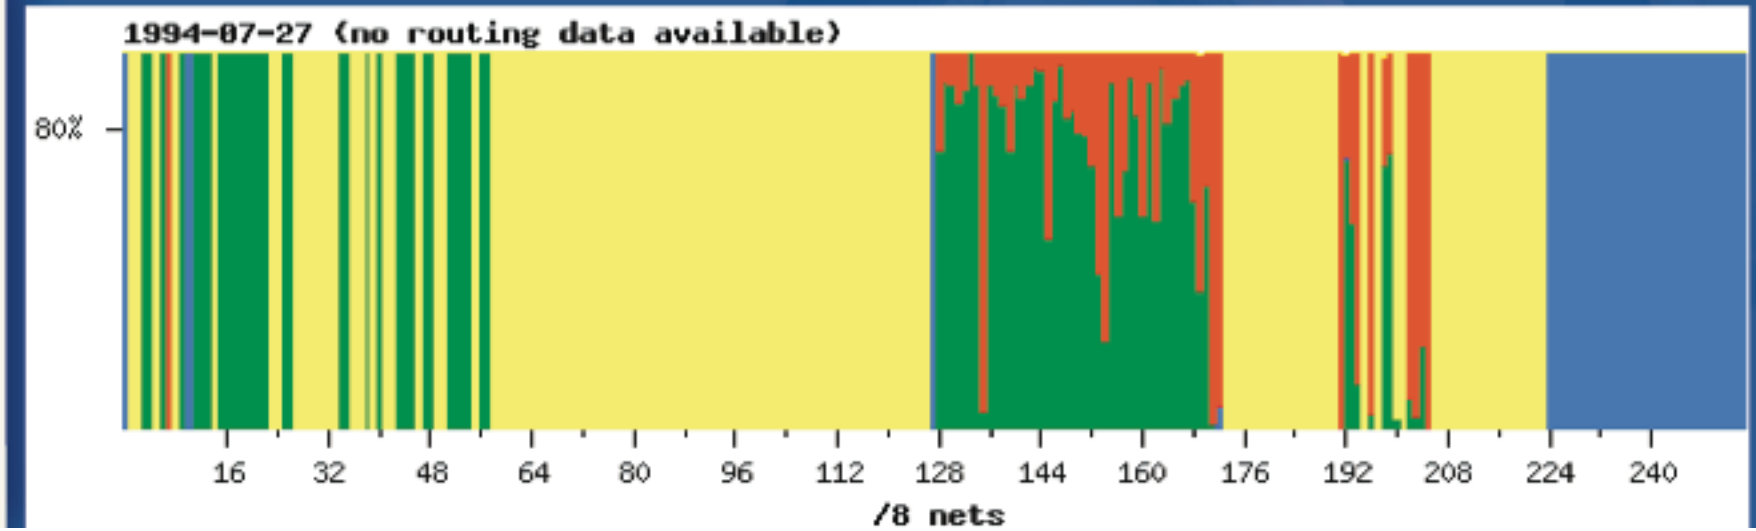

Where did my packet come from?

- Reduce files to ‘most aggregate’ view by economy
  - Published state preserves individual delegations
  - Reduces 100k lines to ~ 35k lines
    - Sorted by prefix/length
    - (de-facto geographic address cloud) *don’t go there…*
    - Include a catch-all ?? Unknown-economy code
      - Darknets and (as yet) untagged assignments
    - Sort data by address
- Apply Simple tape sort/merge algorithm across data, prefix list (Dijkstra, 1978)
  - Brute force, but sort cost amortized over repeat runs through data. Fast to run & re-run.
  - UNIX sort alg highly efficient: sorting data easy(ish)
- IPv4 Address not a good representation for sorting!
    - Convert to %03d padded numerals or HEX to sort

Example: APNIC DNS reverse analysis

- 1 minute tcpdump samples every 15 minutes
- Data from mid-2002 to present
- 7 points of measure (4 primary, 3 secondary)
  - Not all sources present across timeline
- Complements logfile analysis, full packet capture

Queries from JP to AU/HK/JP

- Clear preference for in-country NS
  - High degree of high b/w IXP participation
- AU/HK mostly equal load
  - Slight HK preference? (cable distance?)

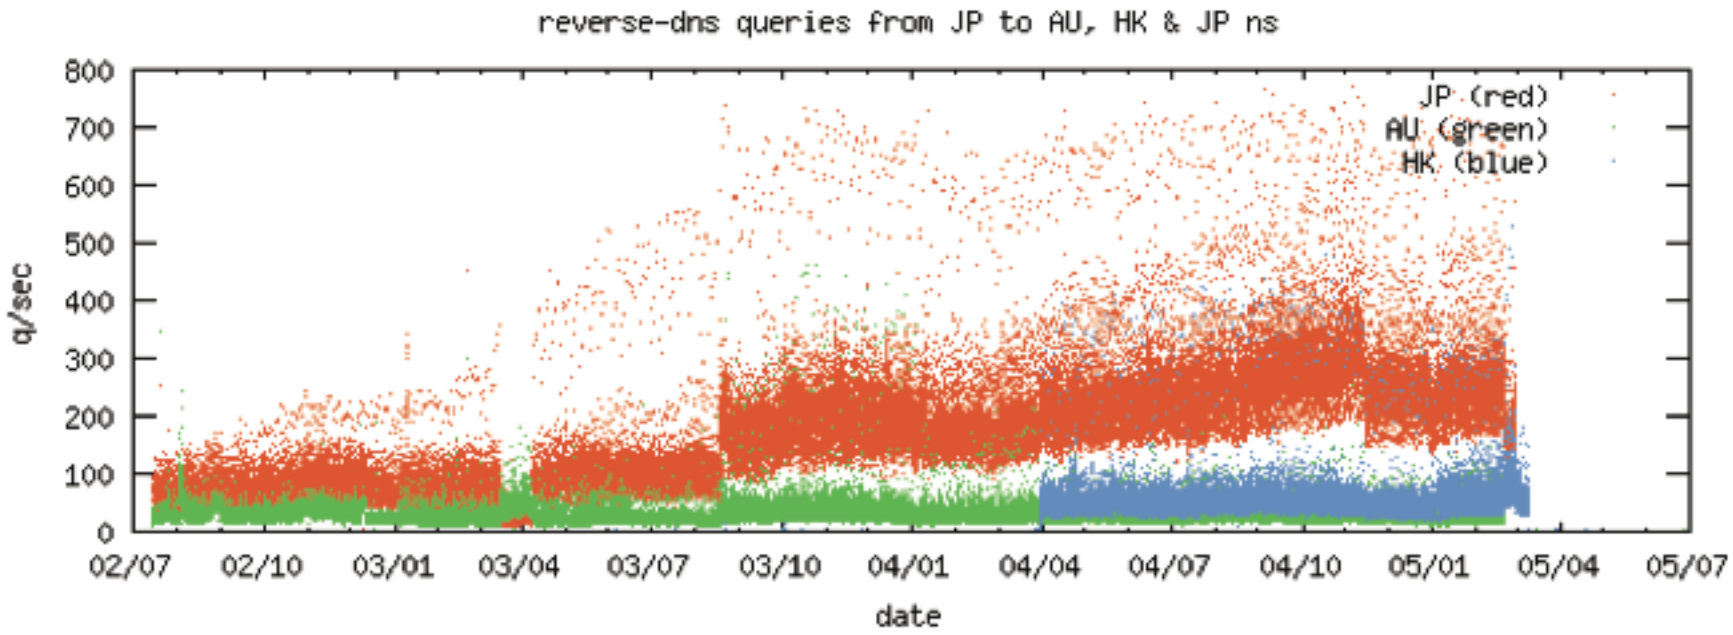

Queries from CN to AU/HK/JP

- Slight preference for JP
  - Until HK node fully commissioned
  - HK seems to take load from both AU/JP
  - Improved IXP participation?

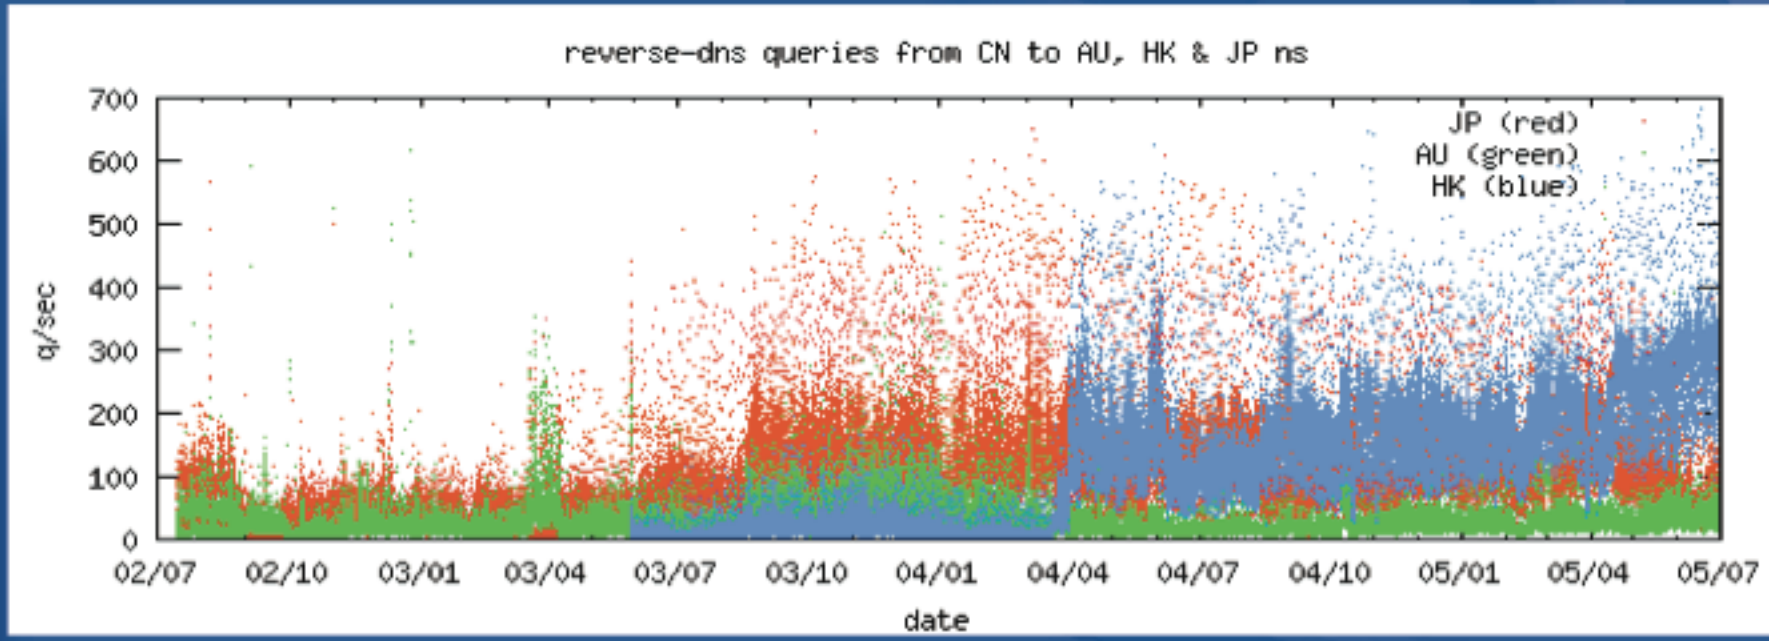

Queries from US to AU/HK/JP

- Slight preference for JP, mostly equal load
- Secondary NS in US/NL serve most
  - (not graphed)
  - Other evidence suggests JP is fast for west coast networks.

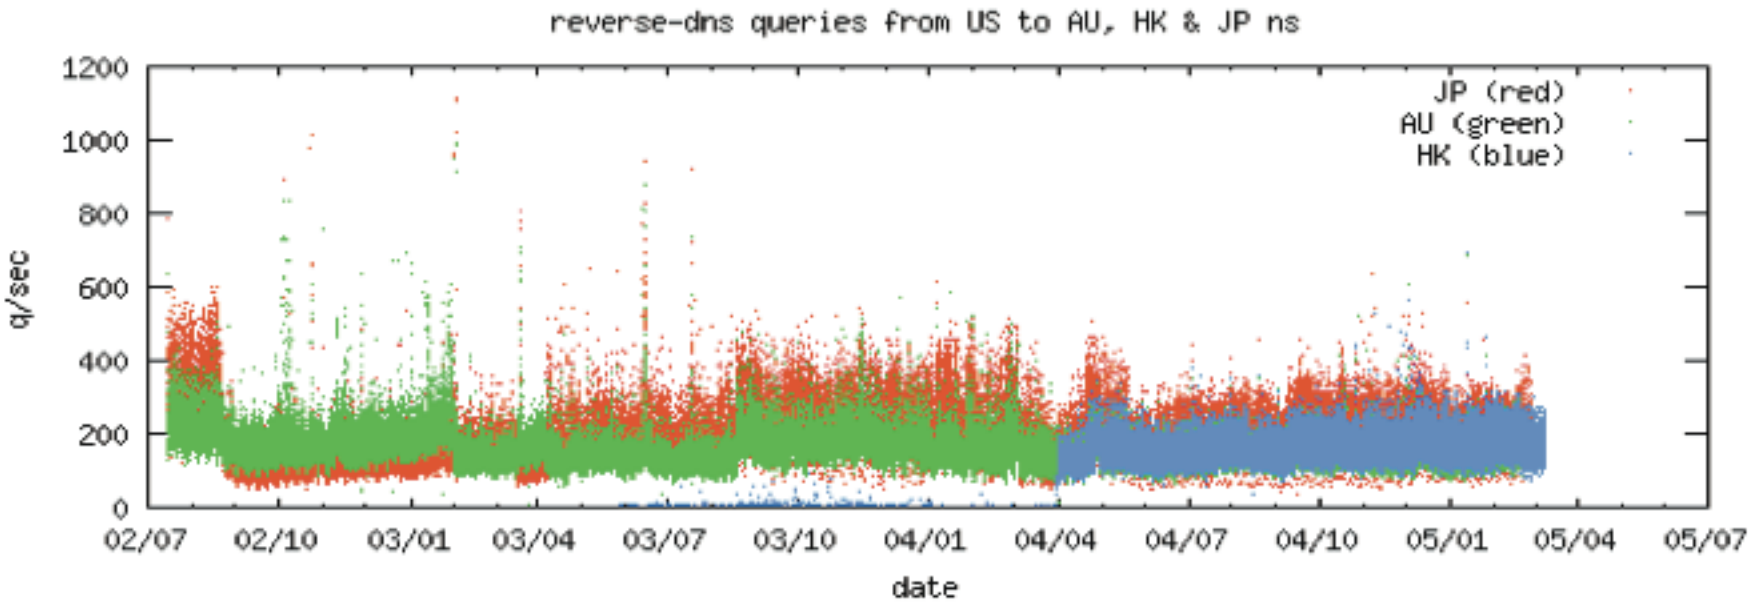

Queries from NZ to AU/HK/JP

- Clear preference for AU
- Some residual traffic to JP/HK

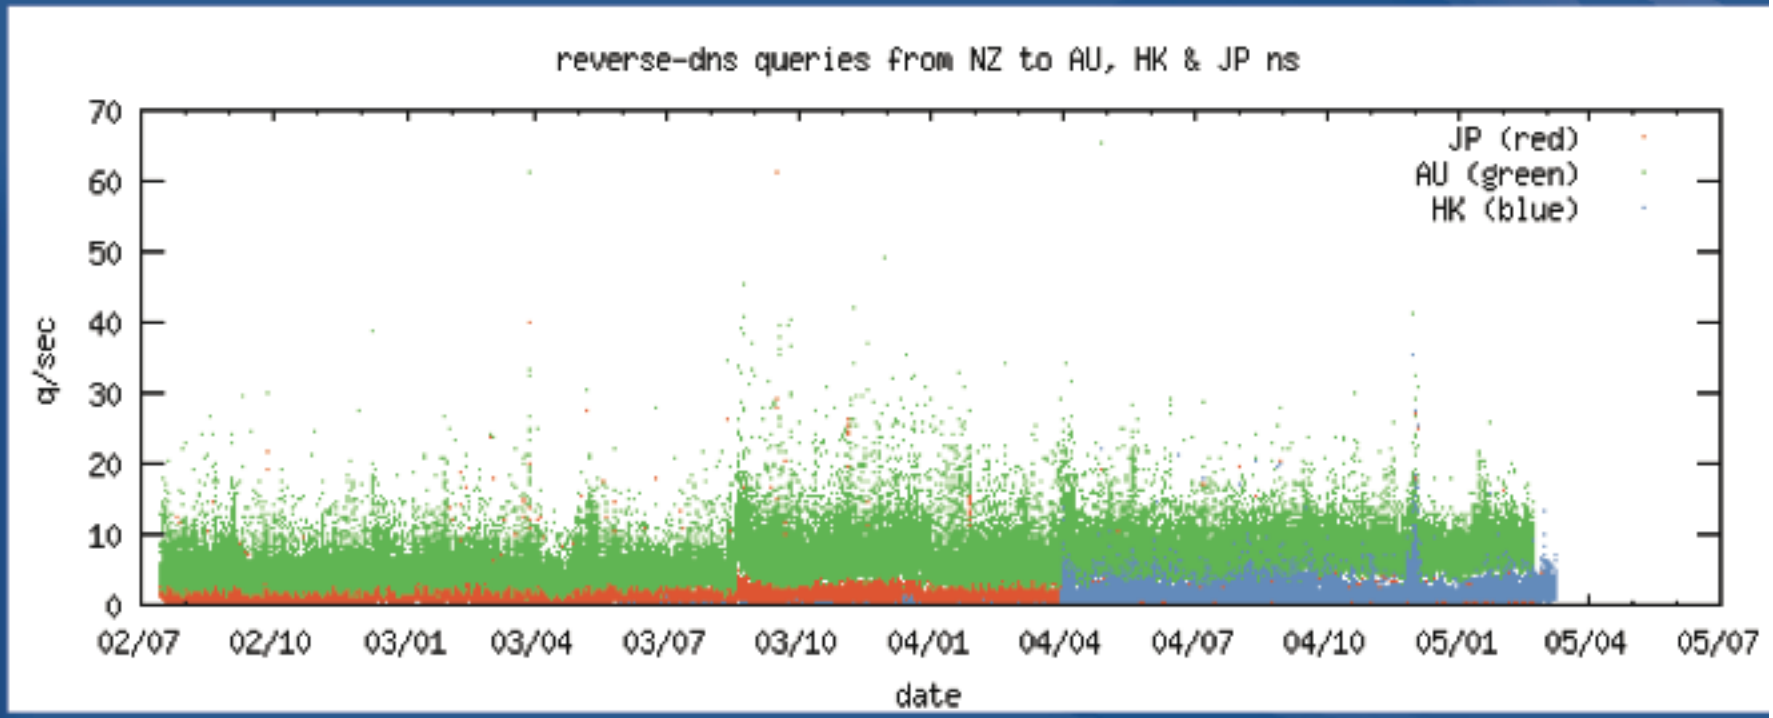

Endogenous vs Exogenous DNS

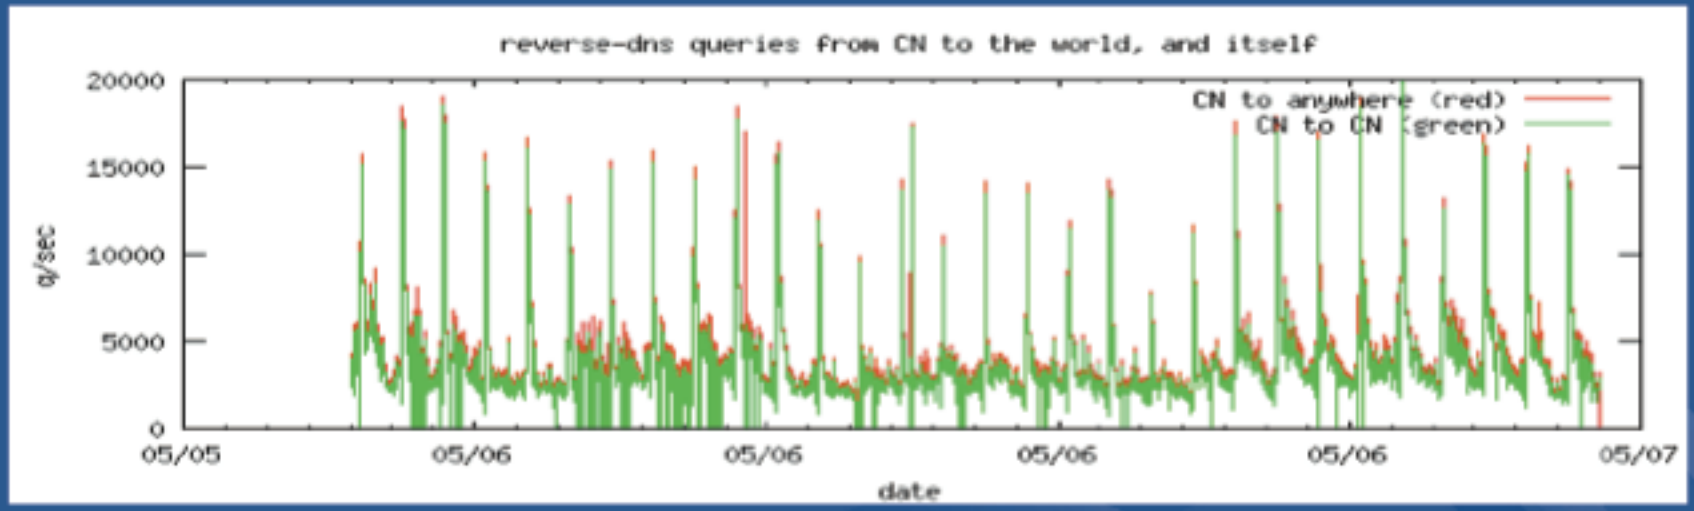

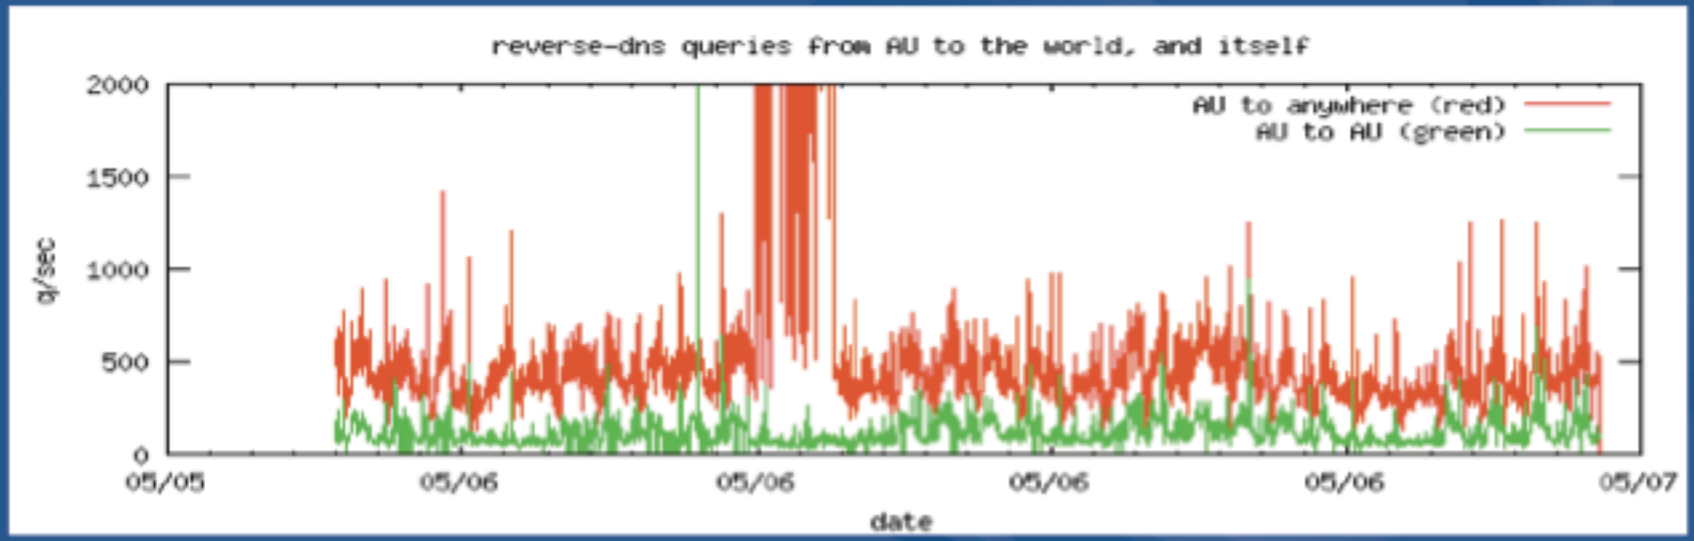

Asia Pacific Network Information Centre

APNIC

Packets-by-Timezone

- Map iso3166 economy to GMT offset
  - Not really applicable to US, CA, RU
- Histogram plot (and animate..)

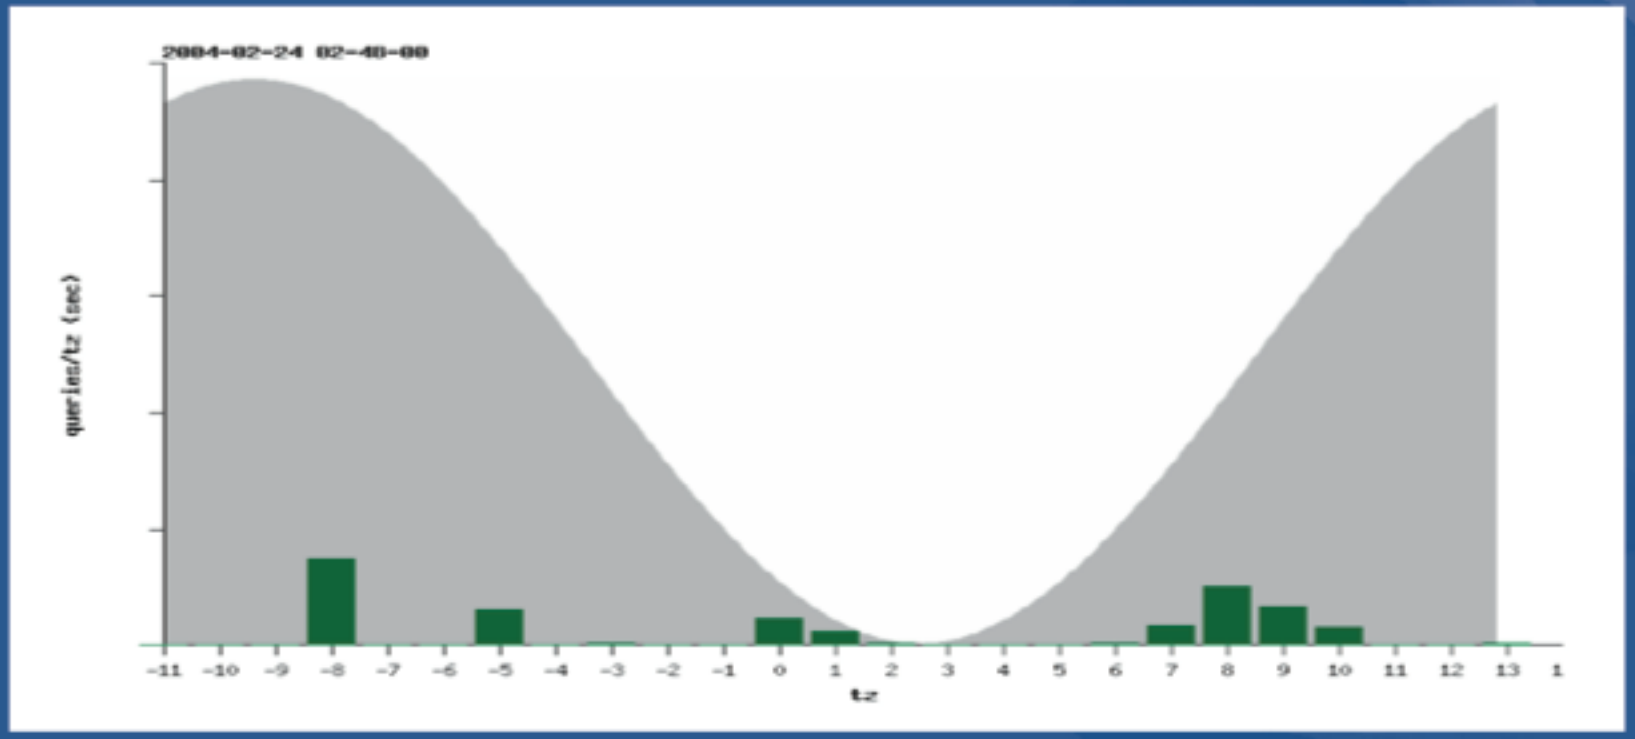

Observations

- Much DNS traffic is in-country
  - JP client looking up JP reverse
  - CN client looking up CN reverse
    - (seen on src, dst plots not shown here)
  - Effect of availability of in-language content?
    - Less extreme for english-language economies
- RTT selection alg seen to work
  - JP client preferences JP located NS
  - NZ client preferences AU located NS
  - CN shows cutover when new (short RTT) service available

Observations #2

- Everyone looks up US reverses..
  - And the US/Europe looks up everyone
- 2 points of NS serve not graphed (yet)
  - Secondaries at ARIN, RIPE-NSS
  - Will shortly own an APNIC US hosted NS..
- Very little ‘out of region’ application of AP addresses
  - Not so true for US/EU delegated resources..
- Same technique applied to Root data
  - (OARC has plots for old H address)
  - Hope to present work on more data ‘rsn’

Proposed Extensions to format

- Two new fields proposed:
  - ‘Same allocation’ marker
    - Using up disjoint space, assign separate elements as one ‘atomic’ event
    - Will help track assignment size behaviours
    - Still not a transaction log
  - ‘Same entity’ marker
    - Helps clarify how much address entities have
    - After M&A many leave pre-existing records untouched
    - Can now tag by ‘real’ owner
    - Can track re-request rate per entity

Customer/Economy Prefix-length data

- JP proposal at AMM Kyoto, to aid with resource consumption planning
  - Remove any identity of entities holding resource
  - Summarize prefix lengths of documented customer assignments by economy
- Proposed extension may adopt same file format with span fields (CSV compatibility)

Things to think about

- ‘Not a transaction log’
  - Can inter-generation checks show rate of transactions?
    - Difference between date of file (change) and date of record in file
  - How to represent hand-backs explicitly?
- Finer-grained economy data missing
  - ‘east coast USA vs west coast’
  - What about IBM GSA and other global entities?
  - Increasing use of non-ISO3166 codes
    - AP, EU (now semi-official)

Things to think about #2

- Cross check with BGP/Economy data
  - Some differences expected
    - Differences may be interesting
  - Existing AS/Economy lists look ‘dirty’
    - Casual checks show inconsistencies
    - Use RIR files as confidence check?
- Smarter processing methods
  - Tree based filters
    - No requirement for sorted input data
    - Fast lookup