APNIC/DNS/DITL

ggm@apnic.net

APNIC

- Regional Internet Registry
  - Delegation for in-addr & ip6 .arpa. Sub-spaces
  - Primary for Asia-Pacific Address ranges
  - Secondary for Europe, Africa, Latin America
- DNS loads of the order 5,000 – 12,000 q/sec
  - 3 points of presence
  - Brisbane, Hong Kong, Tokyo
  - Identical NS deployments each location
    - ANS, single host (at this time)
    - One nserver for primary, one for secondary at each location

APNIC/DNS

- Reverse is a very distinct DNS traffic class
  - – Predictable query, response sizes (PTR record)
    - PTR & NXDOMAIN close to same size (absent DNSSEC)
  - – Typically 1-2 NS listed per reverse range
- All reverse listed under (some of) ~10 NS
  - – RIRs secondary each other
- Partial DNSSEC deployment
    - RIPE NCC only at this time
    - Rest of RIR system expected early in 2010
      - – …which we look forward to capturing in DITL 2010!

APNIC/DNS

- Also some ccTLD serve
    - Very minor traffic set against PTR query load
    - Not discussed in this presentation
- NS set shared with ARIN/RIPE/LacNIC/AfriNIC
    - So DITL measurement is only partial
    - But believed to be ‘representative’
    - Both in-region and out-of-region query sources
- Clean separation primary/secondary load

APNIC/DNS…

- Why do people ‘do’ reverse-DNS?
  - Log processing? (cron, ‘lumpy’)
  - On-demand queries tied to SPAM, firewall
    - Semi continuous. Servers
  - …or are they?
    - We don’t actually know. IP source variance suggests it might be random hosts.
  - Some ‘cultural’ overhangs/differences
    - Very popular in Japan (c/f drafts in DNSOP)
    - Service disruptions when delegation fails
    - Varying levels of delegation across different economies
- RIR committed to providing reverse-DNS
  - Future service models may well increase use
  - Continued rise in query load

APNIC/DNS/monitoring

- On-server since 2002, tcpdump based
  - 15 min cycle time, 1min samples
  - Relocated off box (passive-TAP) in 2008
- DSC since 2008 (passive-TAP)
  - Continuous packet capture, 3(+) day window
  - DITL from same host (just export dnscap files)
- Good correlations between measurements
  - Sampling will probably cease in 2010
  - Useful for rapid deployment, prototyping before passive tap acquired.
  - Not sensible on high query load DNS server

APNIC/DNS/DITL

- Want to understand long term traffic dynamics of our systems
  - Growth/trends (capacity planning)
- Differences inter-site emerge
  - v4/v6?
  - Query load?
  - IX ‘richness’ ?
- What can DNS tell us about address/resource usage?
  - v4/v6 deployment/uptake measurements
  - Informing global policy, OECD/EU measurement processes

Why unique Ips?

- Query load per IP not yet well understood
  - Still trying to define questions!
  - High variances seen in day, between days
- Distinct IP's seem a strong indicator of v4/v6 relativities
  - Consistent with other measurement domains web, routing
  - Allows regional, per-economy, per-entity (ISP) analyses
    - Ips per allocated network measurement of use in resource/policy planning
- Native, Tunneling comparisons possible
  - Possible to see RTT differences?

Typical load patterns (from DSC)

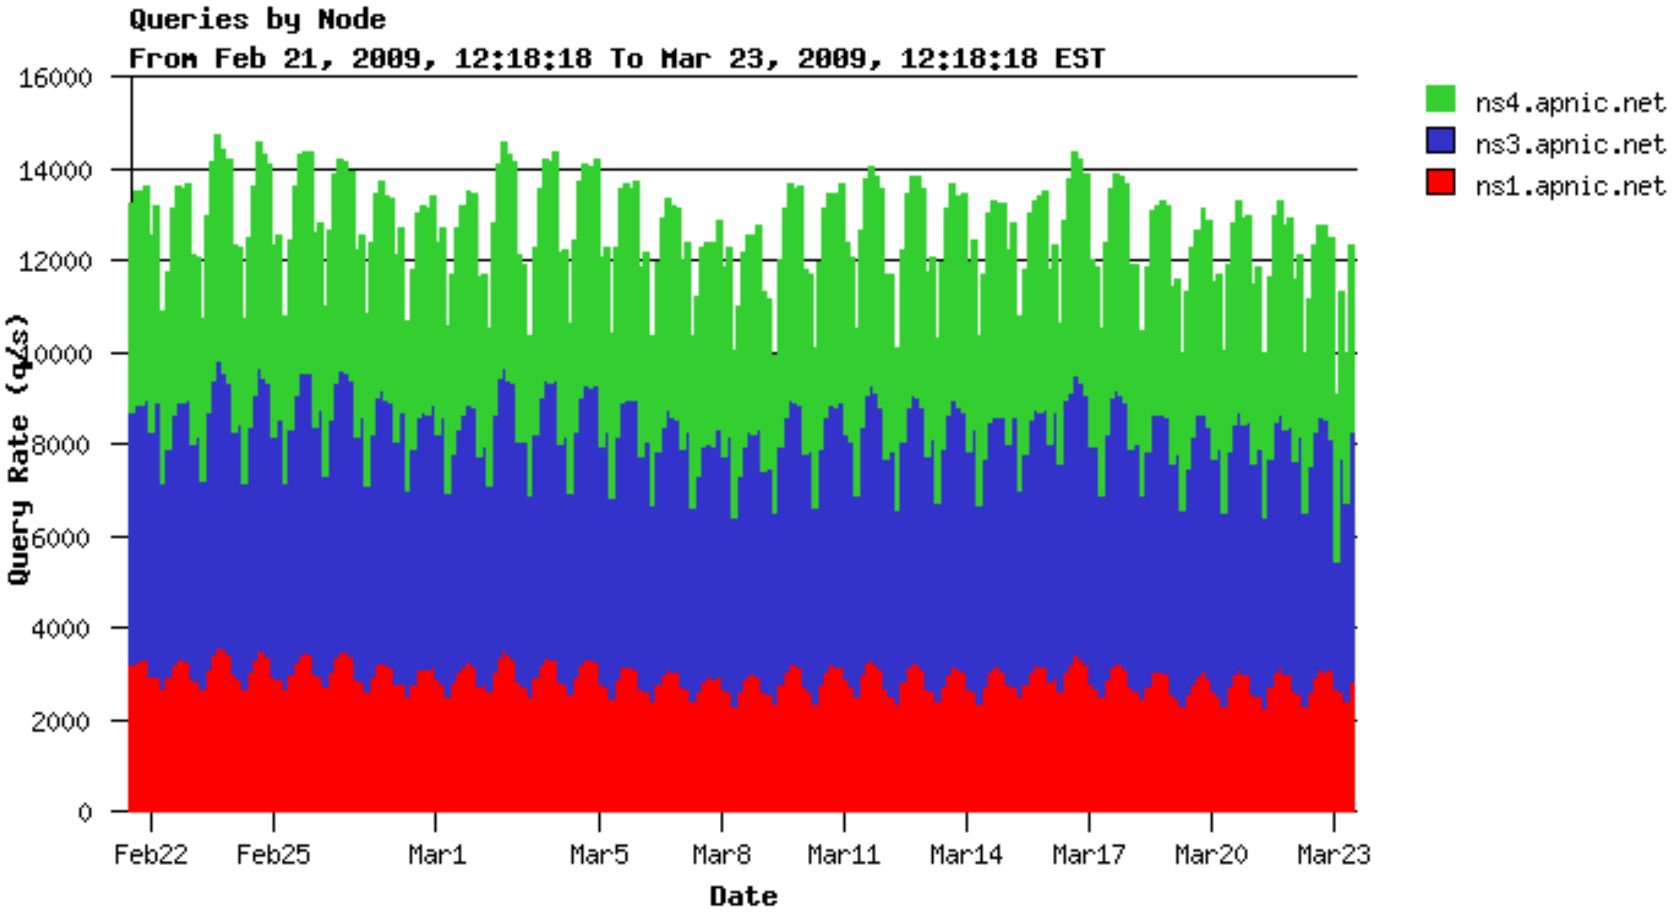

observations

- Strong diurnal pattern
- Small inter-day/weekend variances
- Suggests DITL samples in Mar/Apr good indicators for rest-of-year activity
- Trends/Changes therefore noteworthy
- First approximation: “this process is rational”
  - We’re entitled to treat DITL data as ‘indicative’

Daily traffic shapes (from DSC)

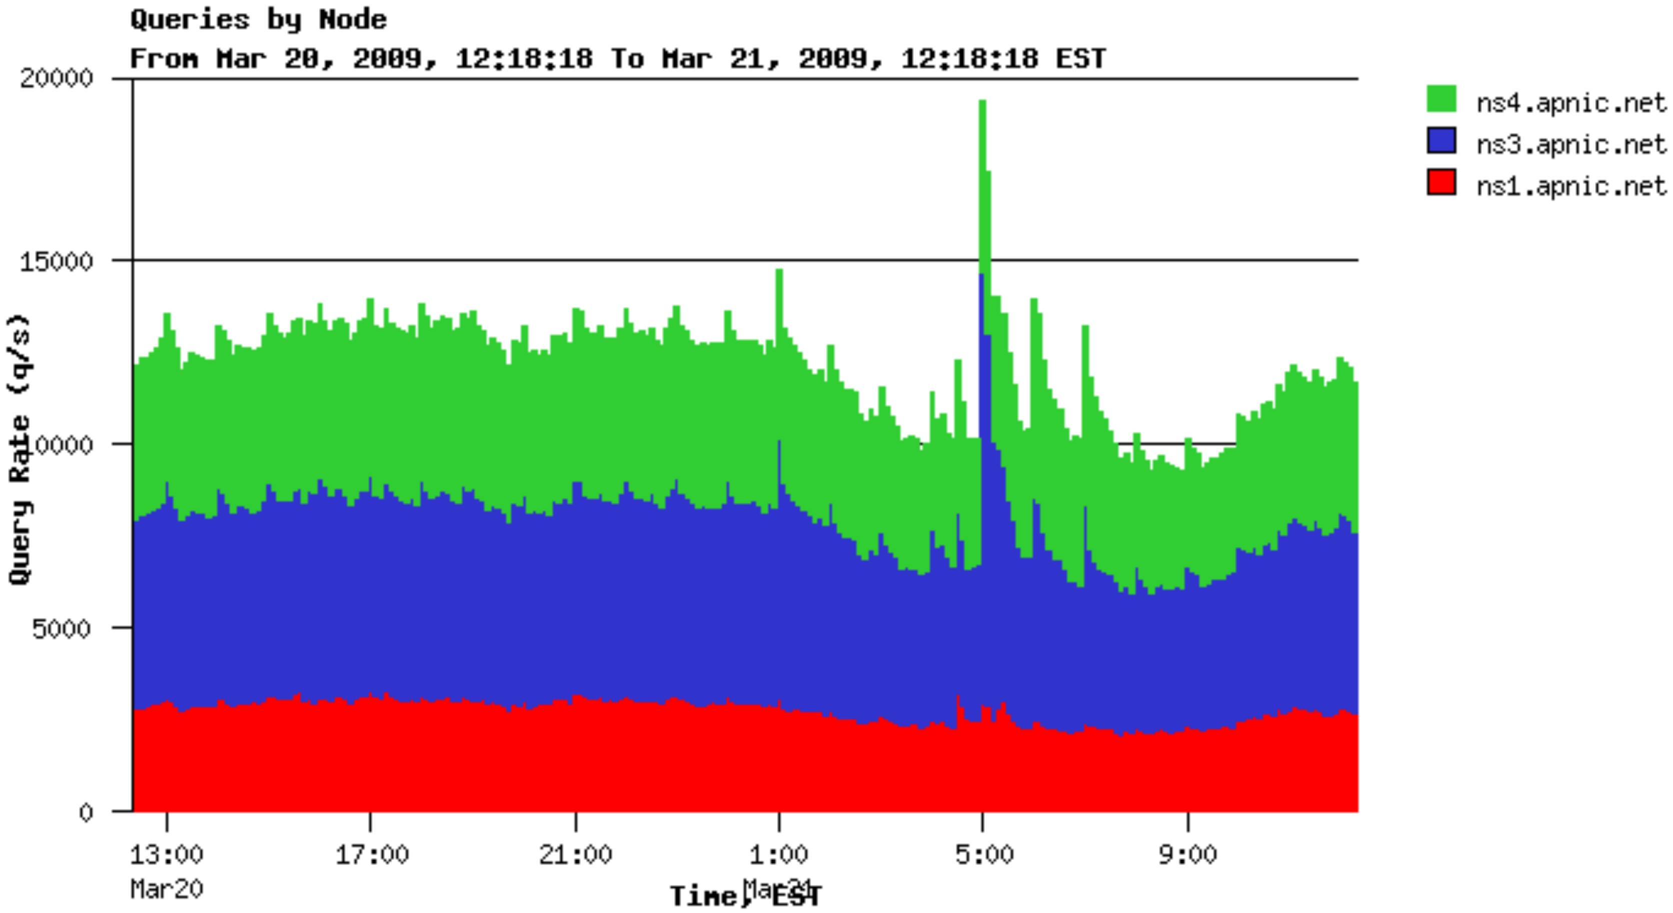

Observations

- Strong single ‘artifact’ required some closer analysis (presented at ESNOG, RIPE)
  - JP ‘signal’ of DNS lookup
- Other ‘clock tick’ interval peaks visible
  - Against background of constant(ish) load
  - Less ‘cron’ affected than we thought.
- Inter-site differences
  - Clearly volume, but also some of the artifacts
  - Strong unity of diurnal pattern for the NS (primary) server set, cross-site.

Unique IPs in 24h

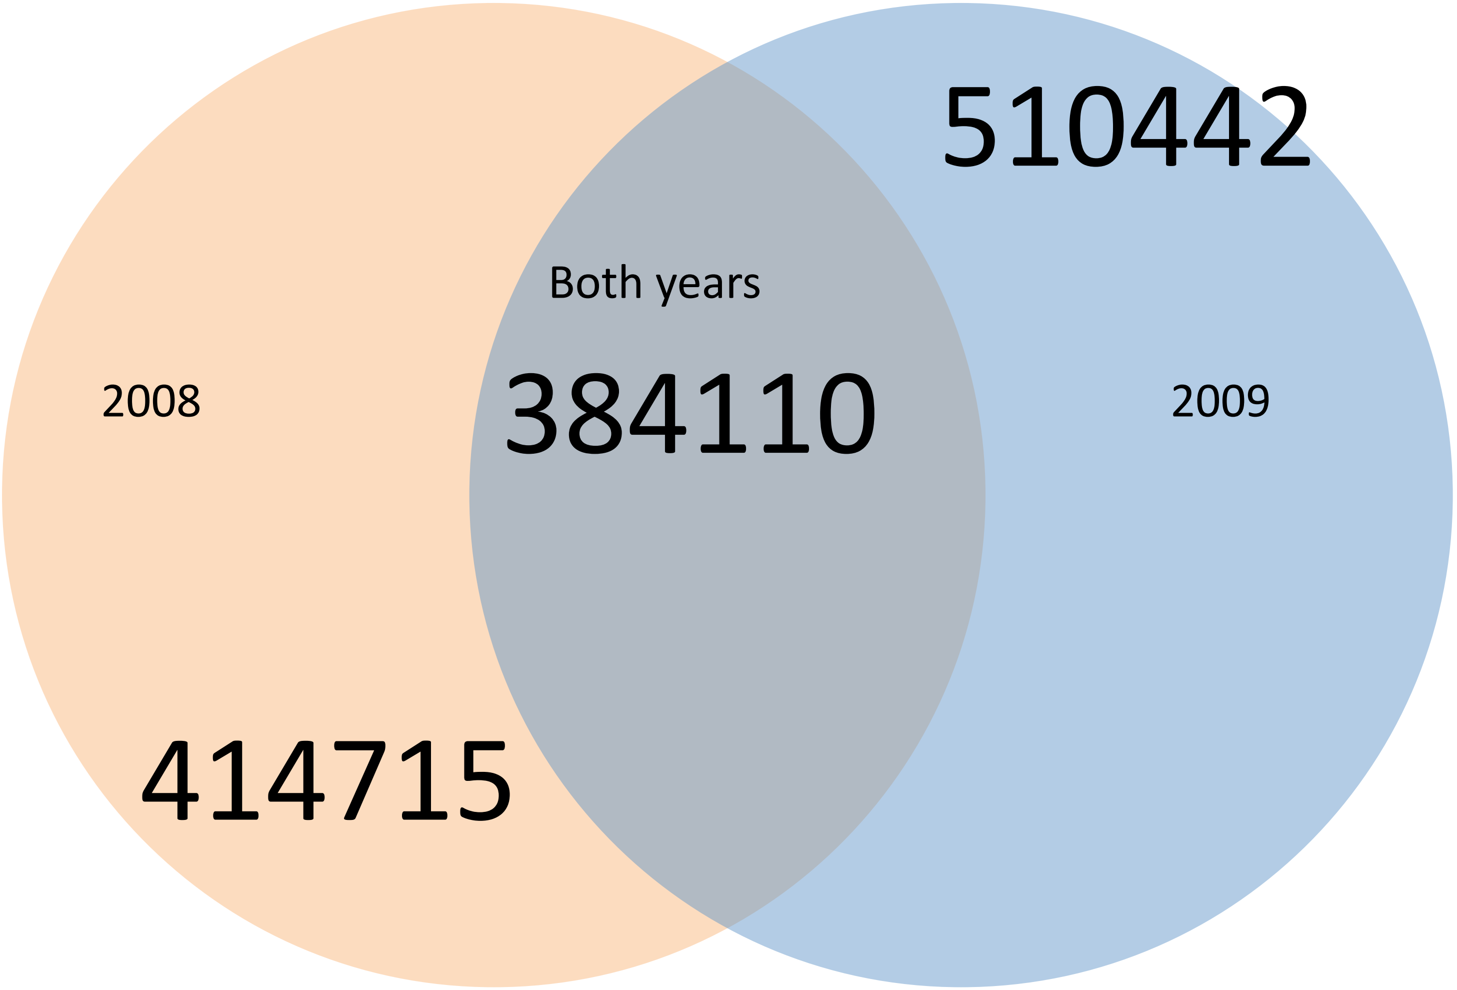

How often do people query?

It's a LOG Scale.

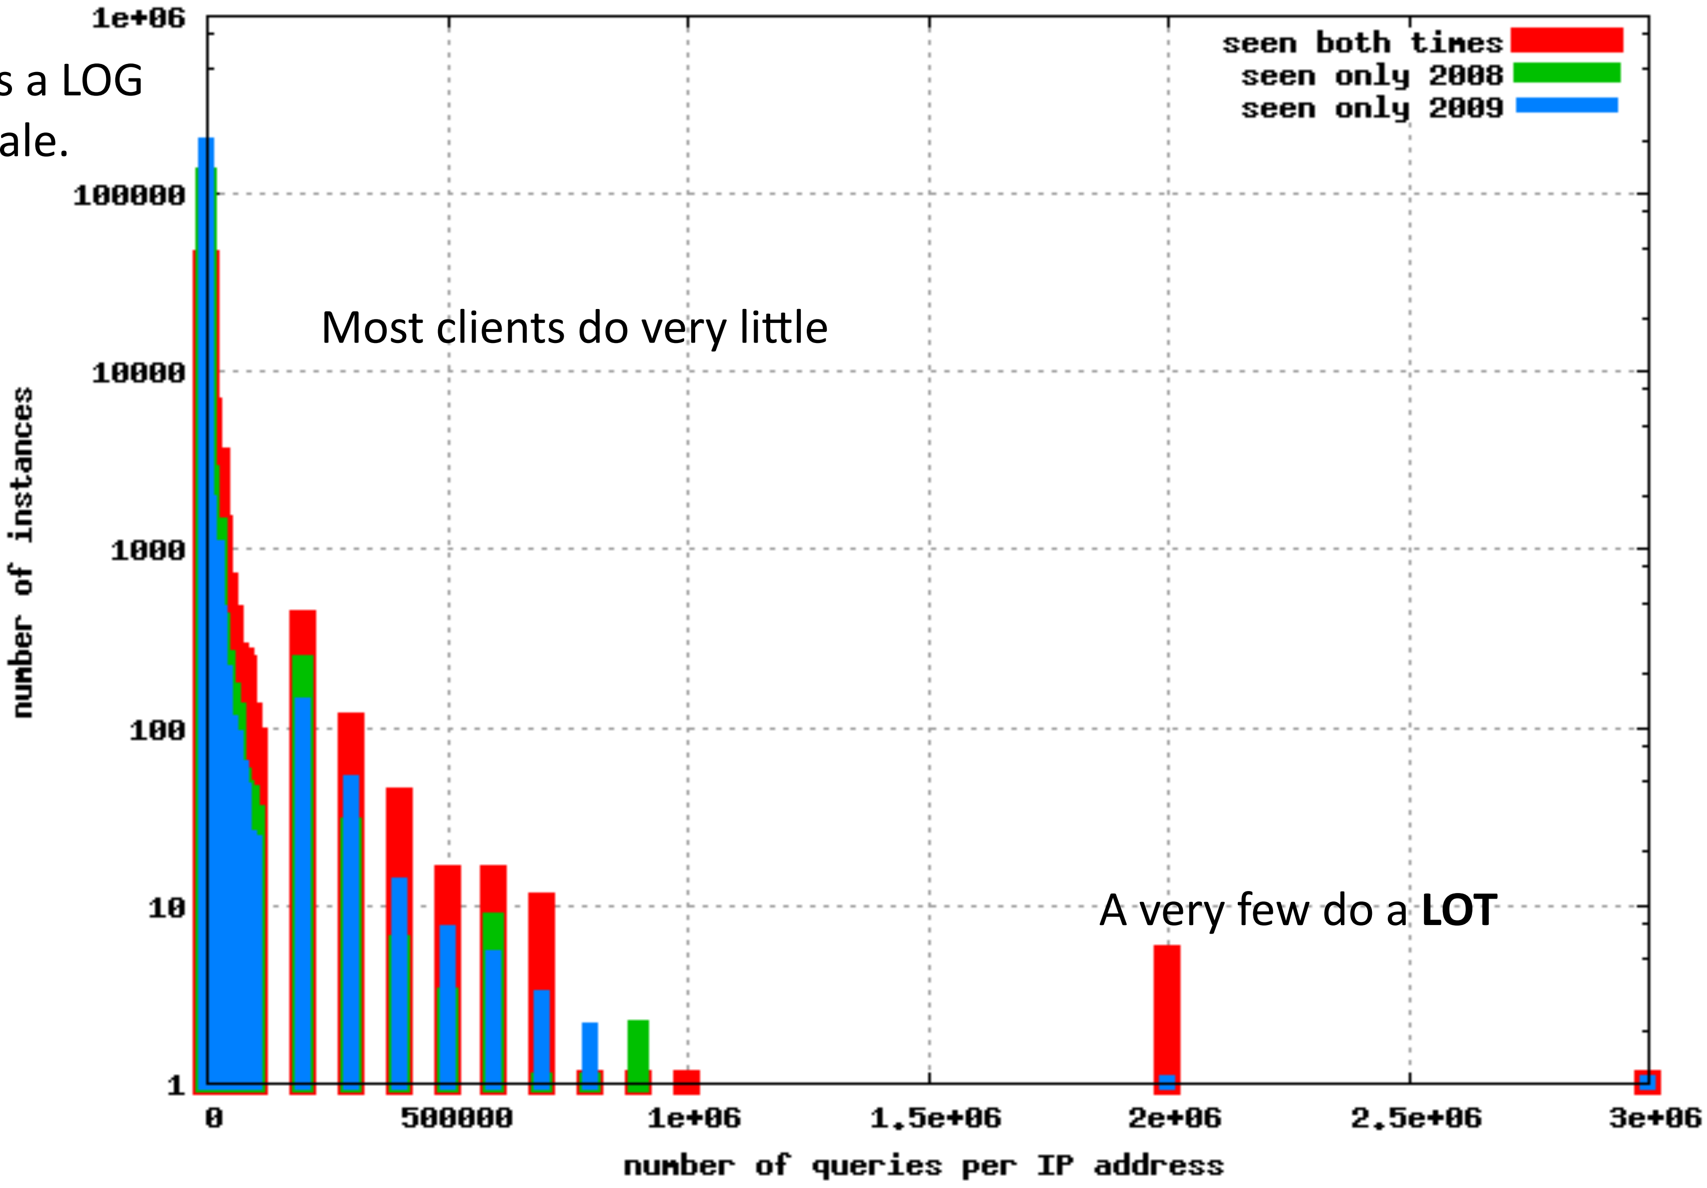

Most clients do very little

A very few do a **LOT**

Client behaviors

- More variance in query pool than expected
- Many visits from Ips that only do 1-2 lookups
- Little address persistance 2008/09
  - – Want to see how this maps out in DITL 2010
- Beginning to analyze by RIR allocation records
  - – Up to 70% of deployed nets ‘seen’ in 2009 DITL
  - – Also want to track this into DITL 2010

2008 V6/V4 ratios

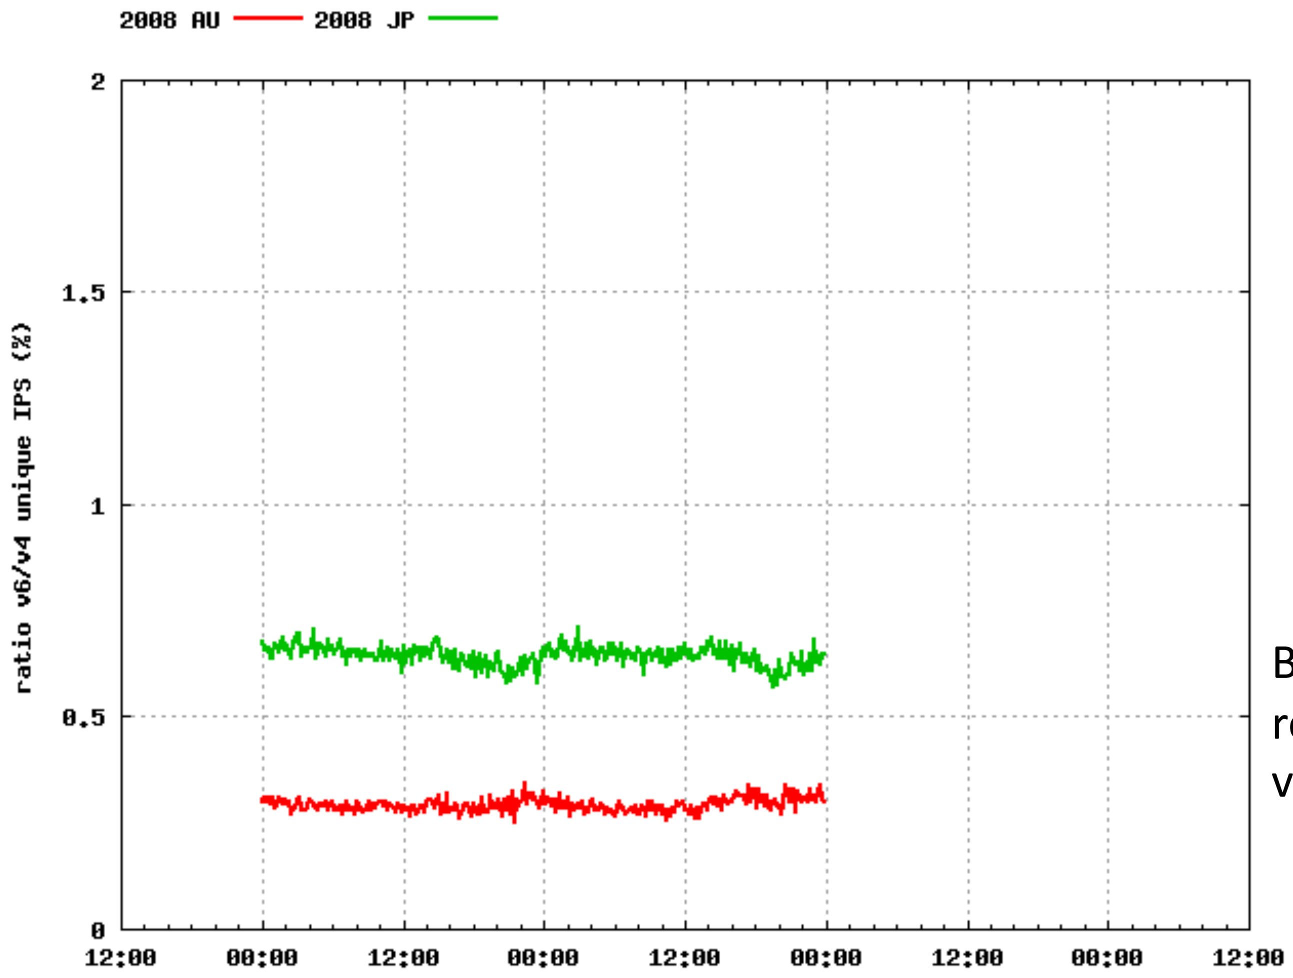

Below 1%, relatively large variance JP/AU

2009 V6/V4 ratios

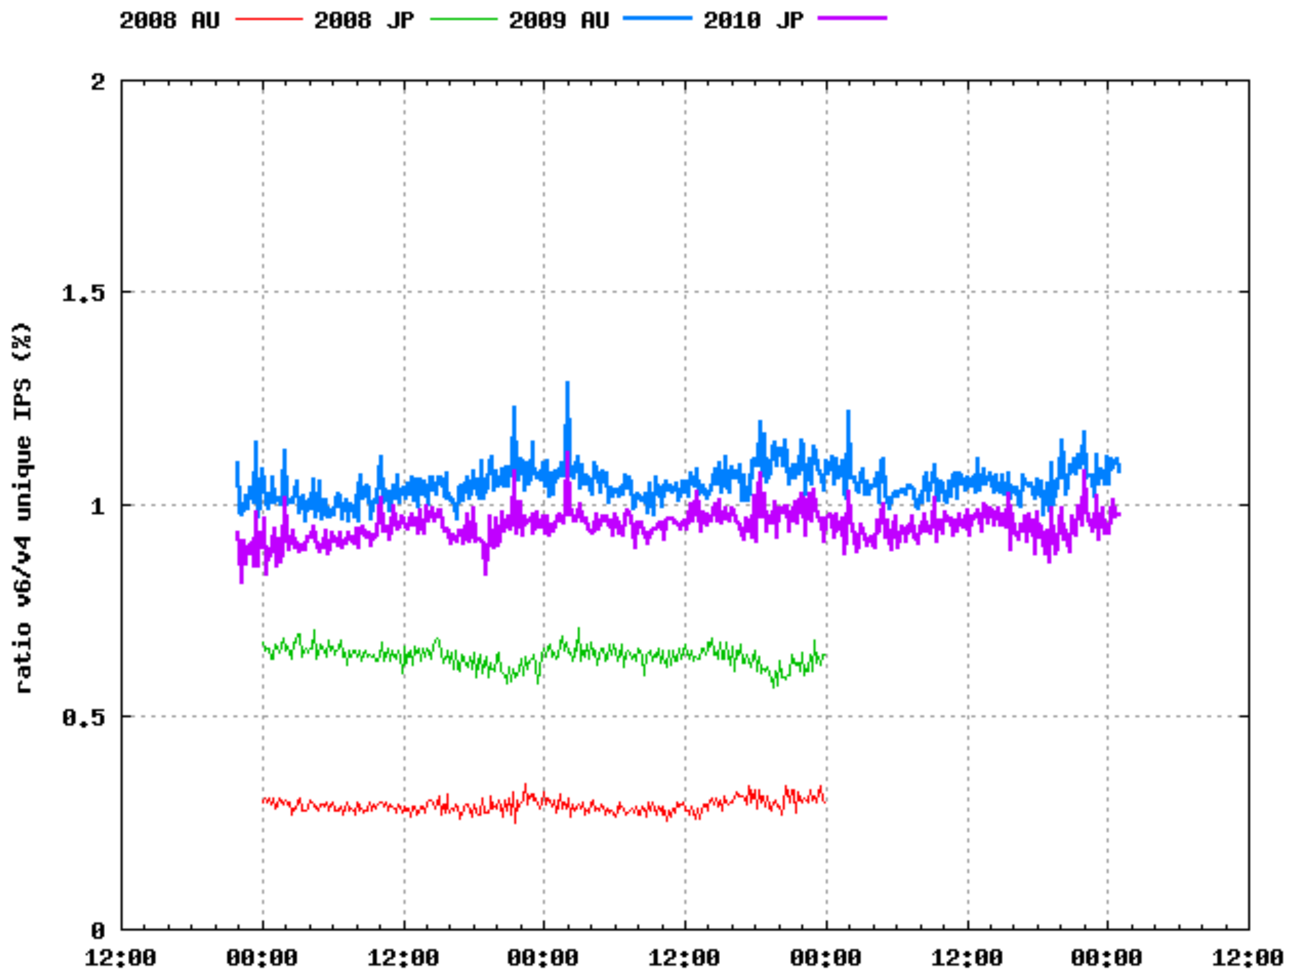

Around 1%,
less variance

Observations

- Weak diurnal pattern: ratio of v4/v6 is not affected strongly by time-of-day
- ‘increased’ use of IPv6 transport in 2009 compared to 2008.
- Also less variance in v4/v6 suggesting improved transit (AU is on tunnel, added peering to HE)
  - Will re-home V6 onto native feed before 2010
- 1% figure is consistent with other V6 uptake measurements for 2009
  - Reverse DNS appears to be good indicator. Useful!
  - May be worth mapping backwards into other measurments
- Worth looking at 2010 for comparitive/trends

V6 tunnels in DNS AP NS

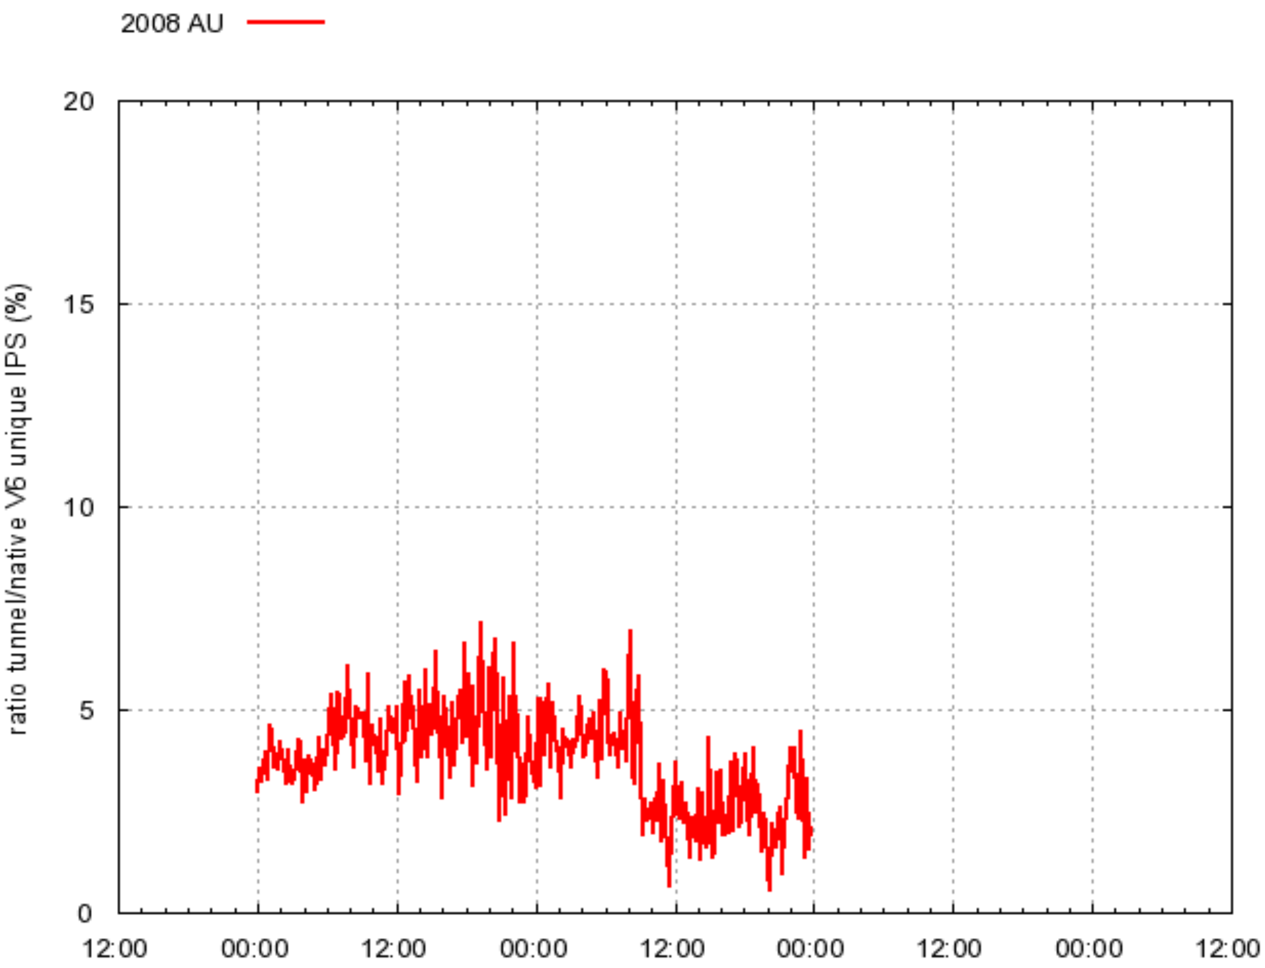

V6 tunnels in DNS AP NS

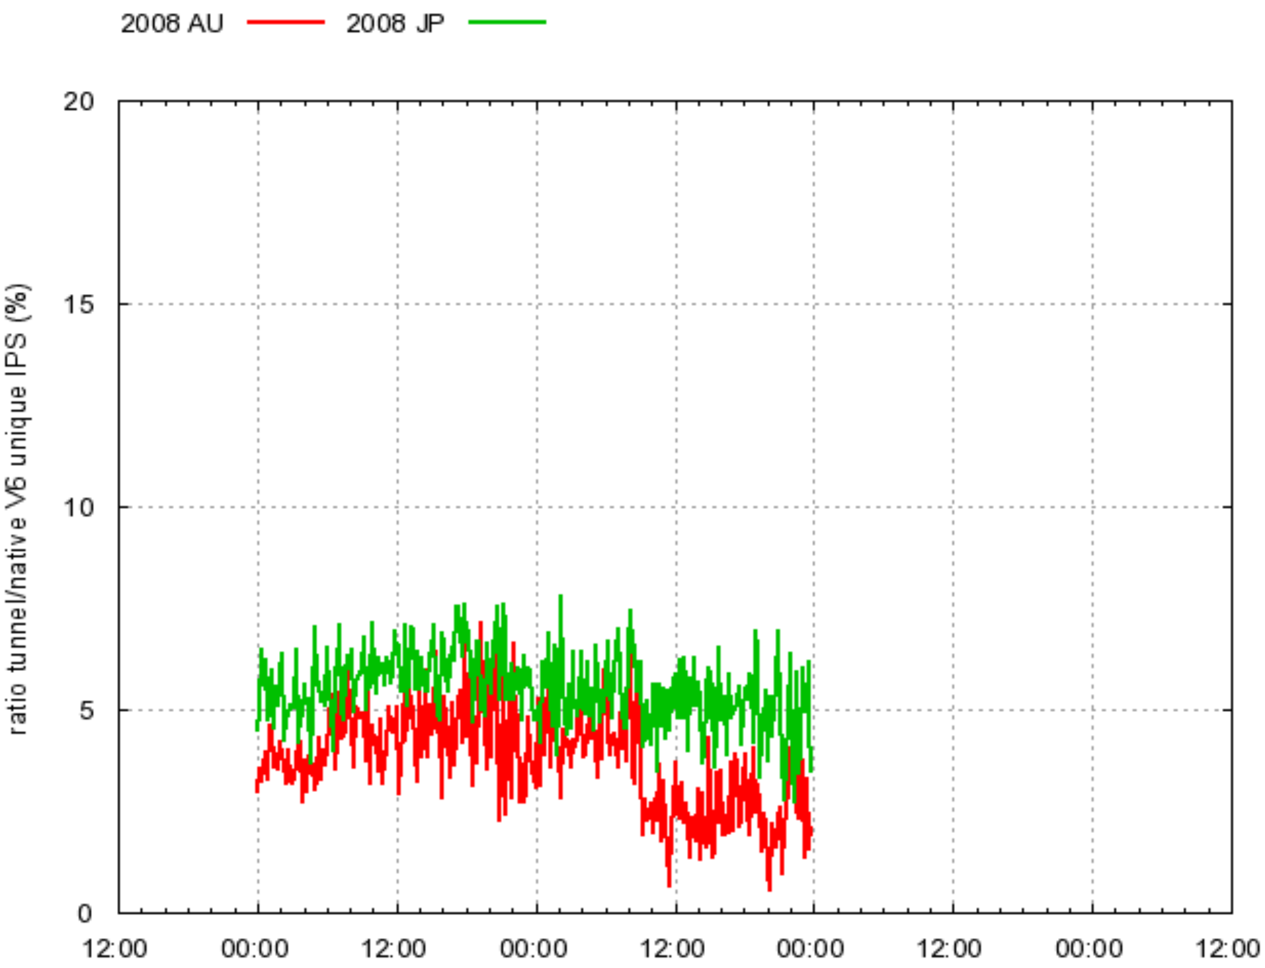

Much more
Variance
per sample

V6 tunnels in DNS AP NS

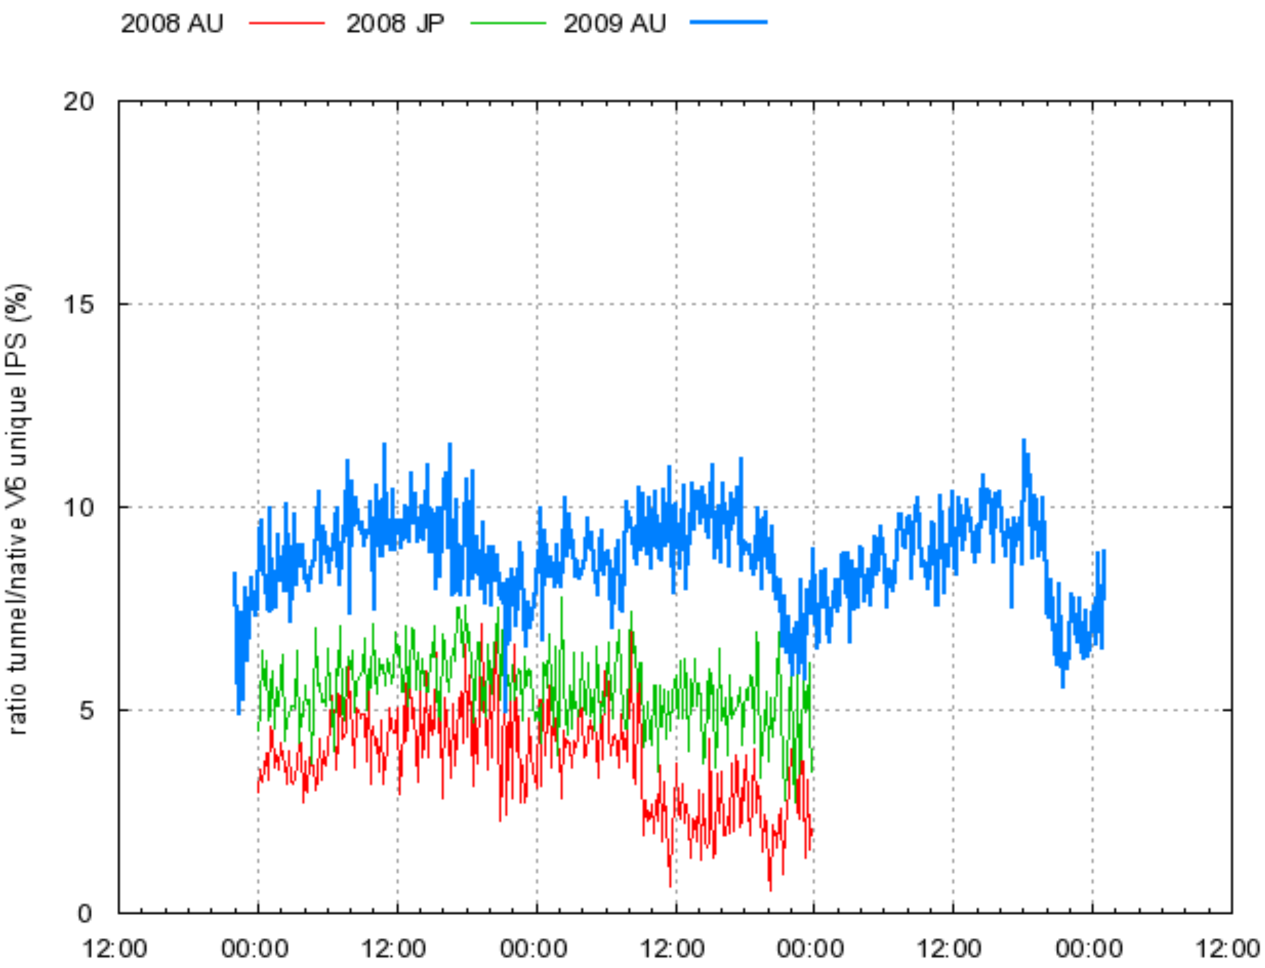

V6 tunnels in DNS AP NS

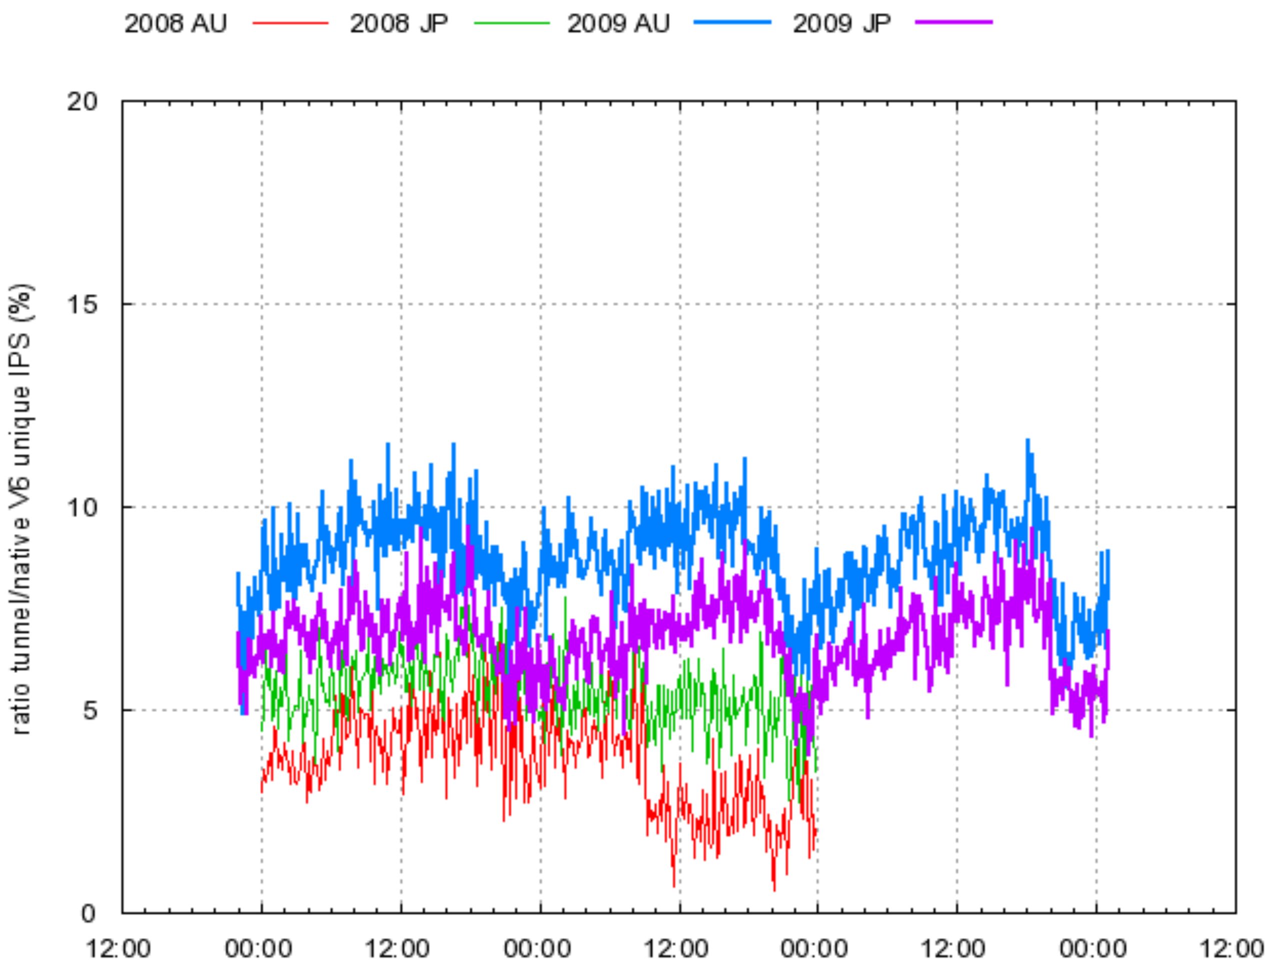

Increase,
Stronger
Diurnal

Observations

- Tunnel dynamics quite different to native v4/v6 ratios
  - Diurnal patterns visible (strongly in RoW servers)
- Some increase in tunnel usage 2008-09
  - Probably reflects HE but may be more widespread
  - Interesting to monitor in DITL 2010
- Clearly sufficient DNS now flows on tunnels that the RTT is ‘viable’ for some people
  - Can detect persistent use of Teredo, 6to4, 6rd
- Native/Tunnel ratios being studied in DNS, web
  - Same scripts/methodology

V6 tunnels in DNS RoW NS

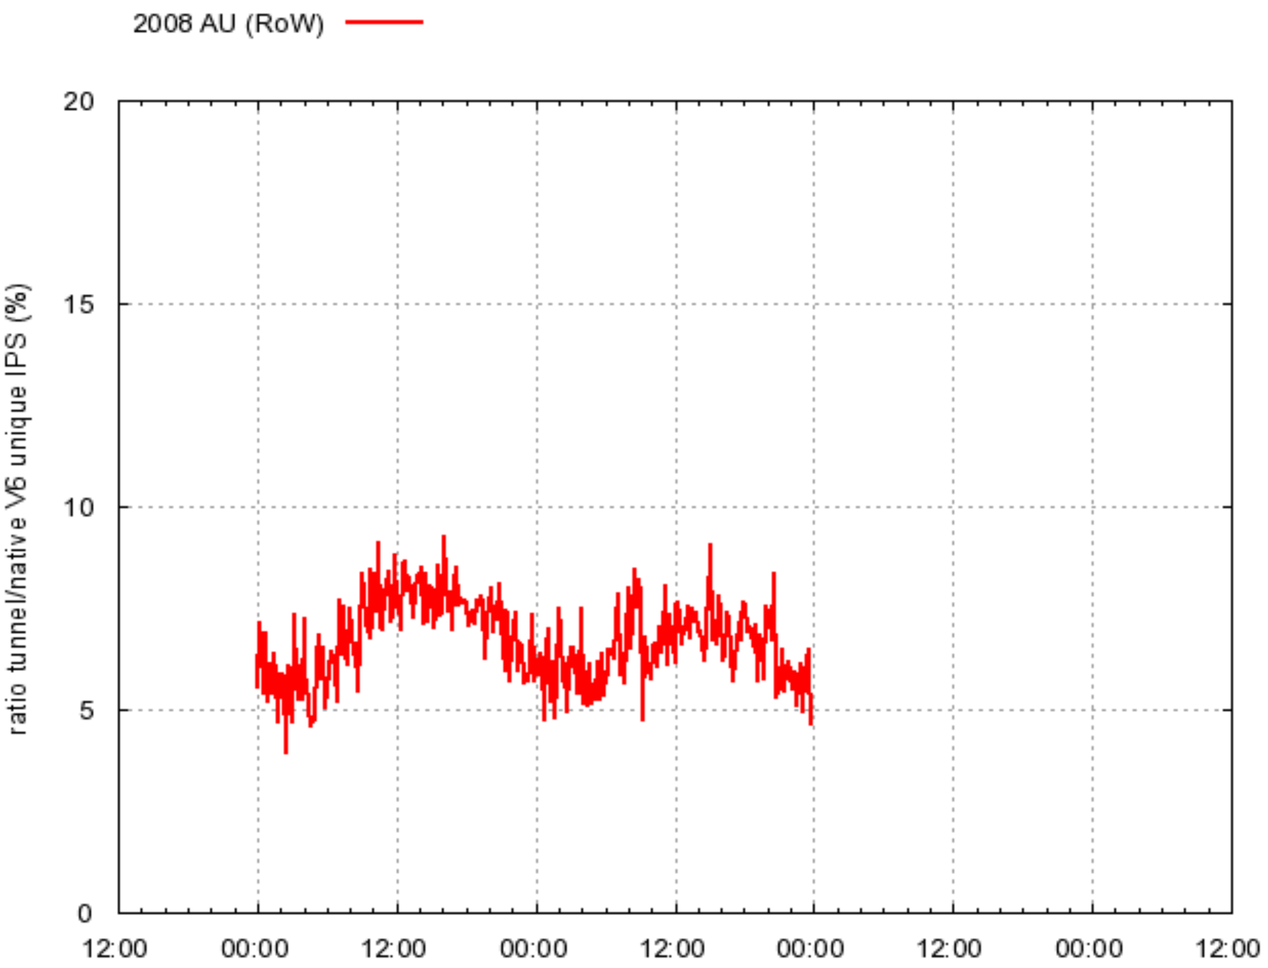

V6 tunnels in DNS RoW NS

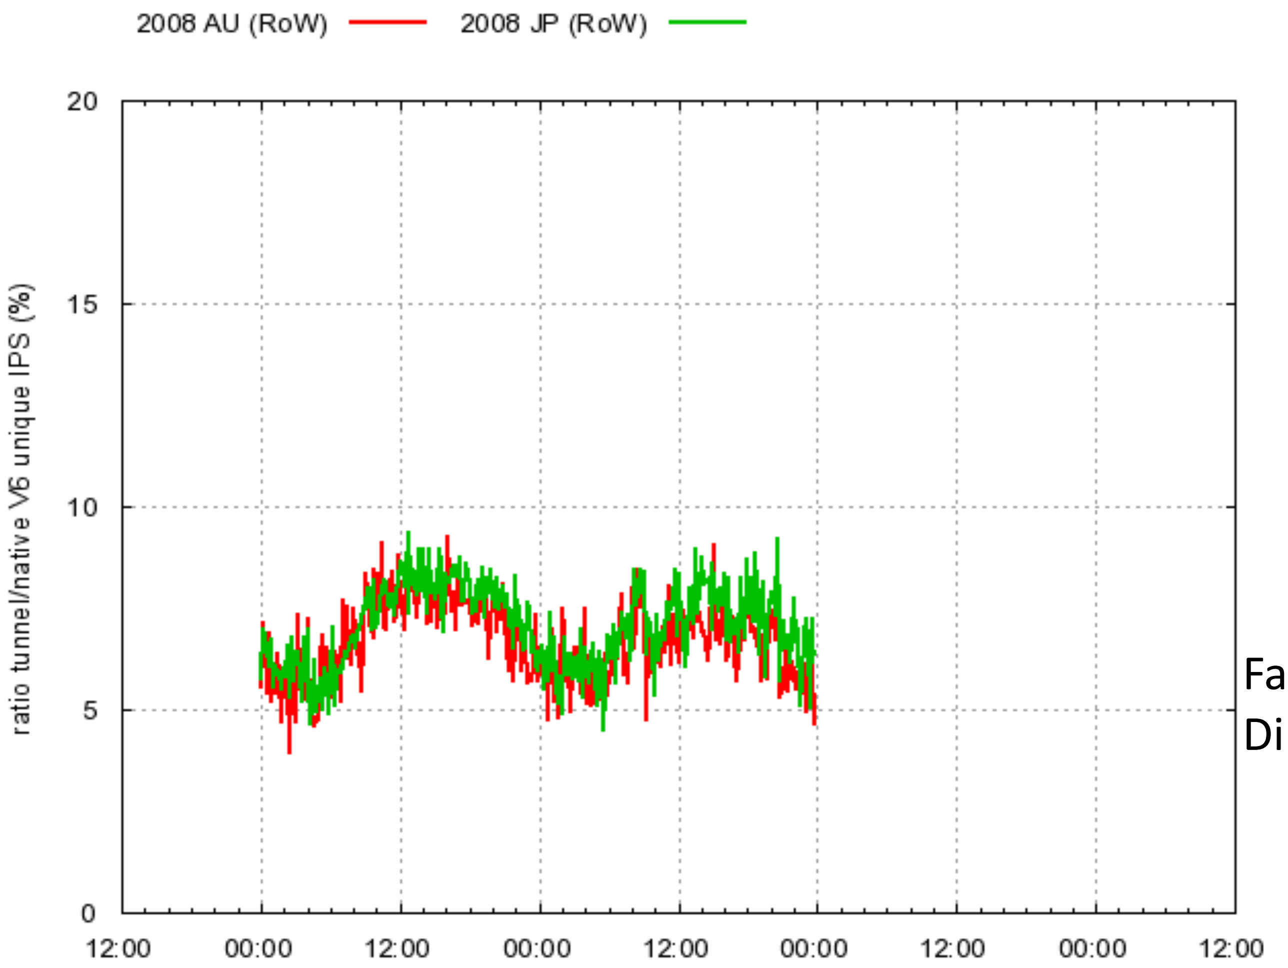

Far less Variance,
Diurnal Pattern

V6 tunnels in DNS RoW NS

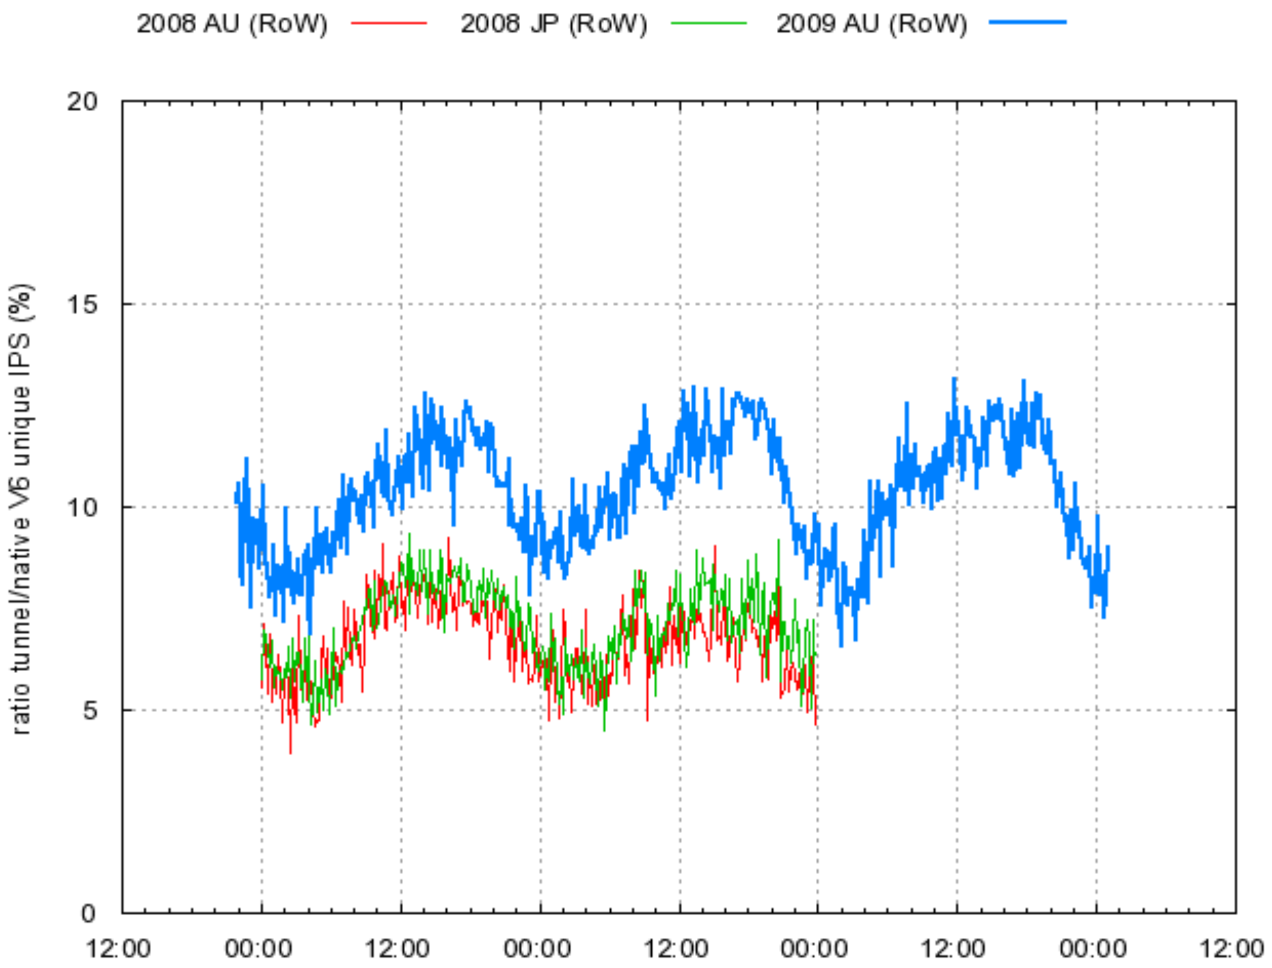

V6 tunnels in DNS RoW NS

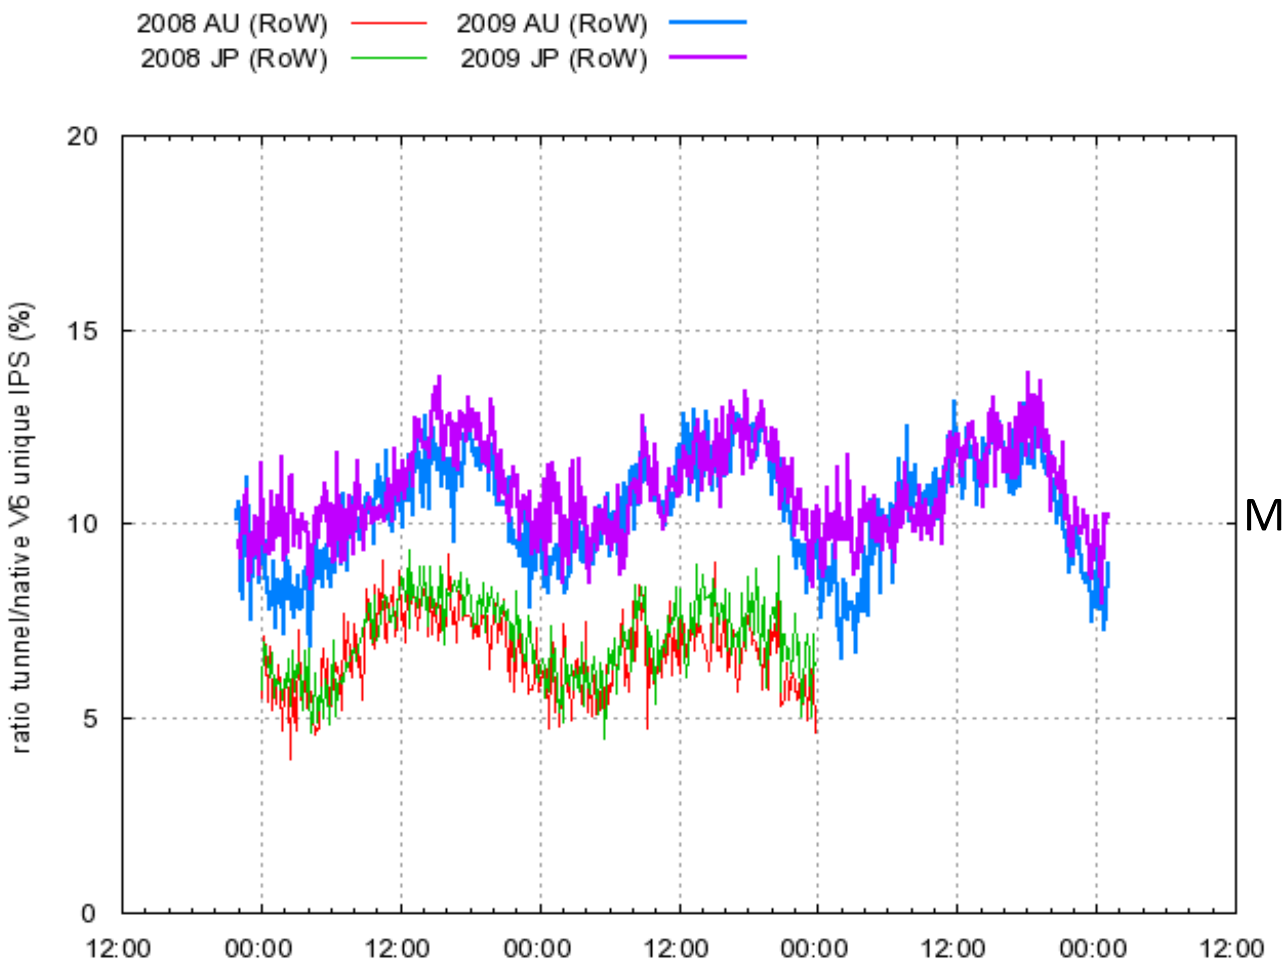

Marked upswing

Observations

- Far ‘tighter’ pattern of ratio for two distinct locations
  - Rest of World includes a lot of out-of-region query load
    - Therefore tunnel vs native V6 RTT probably reduced impact: traffic equally awful (!)
  - Would be interesting to compare with data from the RIPE/LacNIC/ARIN/AfriNIC regions

Inter-economy comparisons

- Take source, destination (PTR) Ips
- Map to economy using RIR ‘delegated’ files
- 2D matrix of {src,dst} pairs 250x250
- 3D (time series) one 2D record per period
- What emerges?
- Animated plots of {from, to} by economy

US looking everywhere (we don’t secondary US reverse)

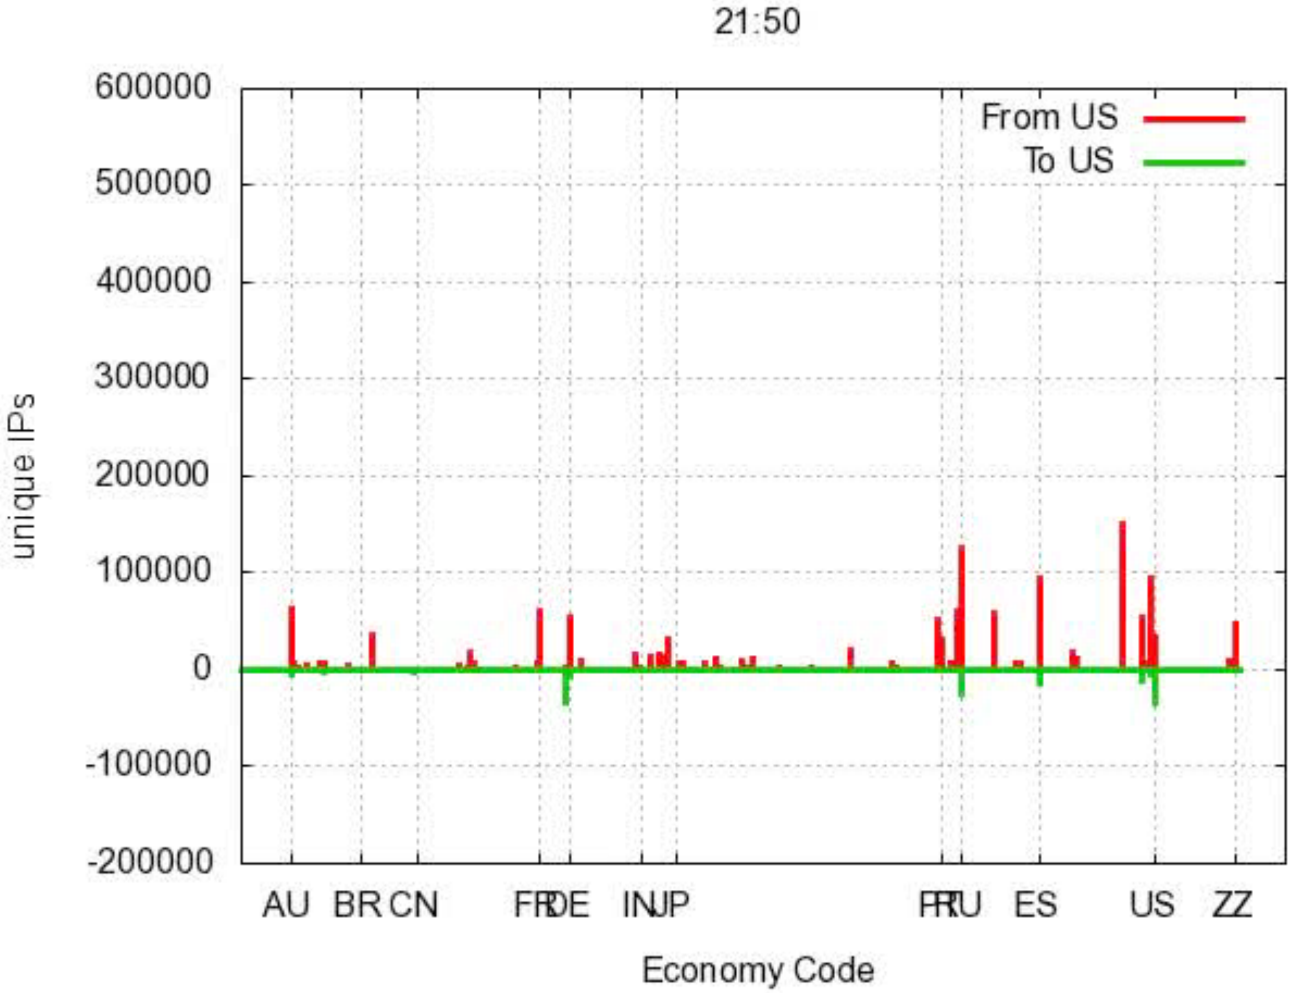

Observations

- Small number of ‘US’ tagged ranges maintained in APNIC reverse maps
- Confidence check #1
  - since we don’t secondary US ranges, did not expect to see large volume of queries with ARIN region *dst* address
- Confidence check #2
  - US economy *src* addresses seen making significant volumes of query to all *dst* addresses

Japan looks at itself

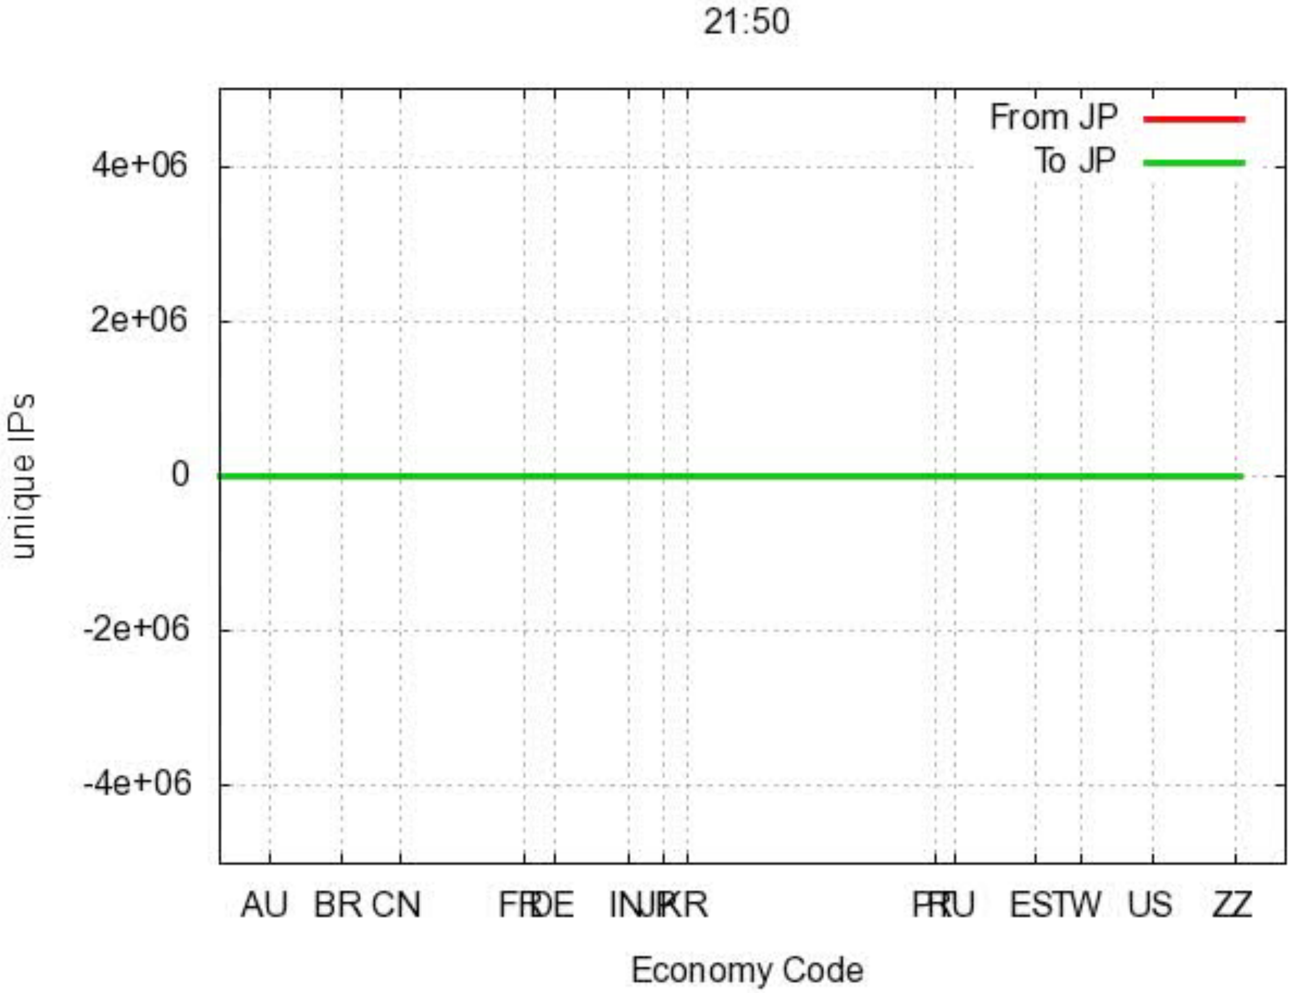

Observations

- {JP,JP} src/dst pairs swamp all other data for JP ranges (notice US over far rhs for scale)
- Suggests intensely internal/domestic DNS activity
- Strong peak. Single measure? Few queriers?
- For further study..

Ok. Japan looks everywhere (but mainly at itself)

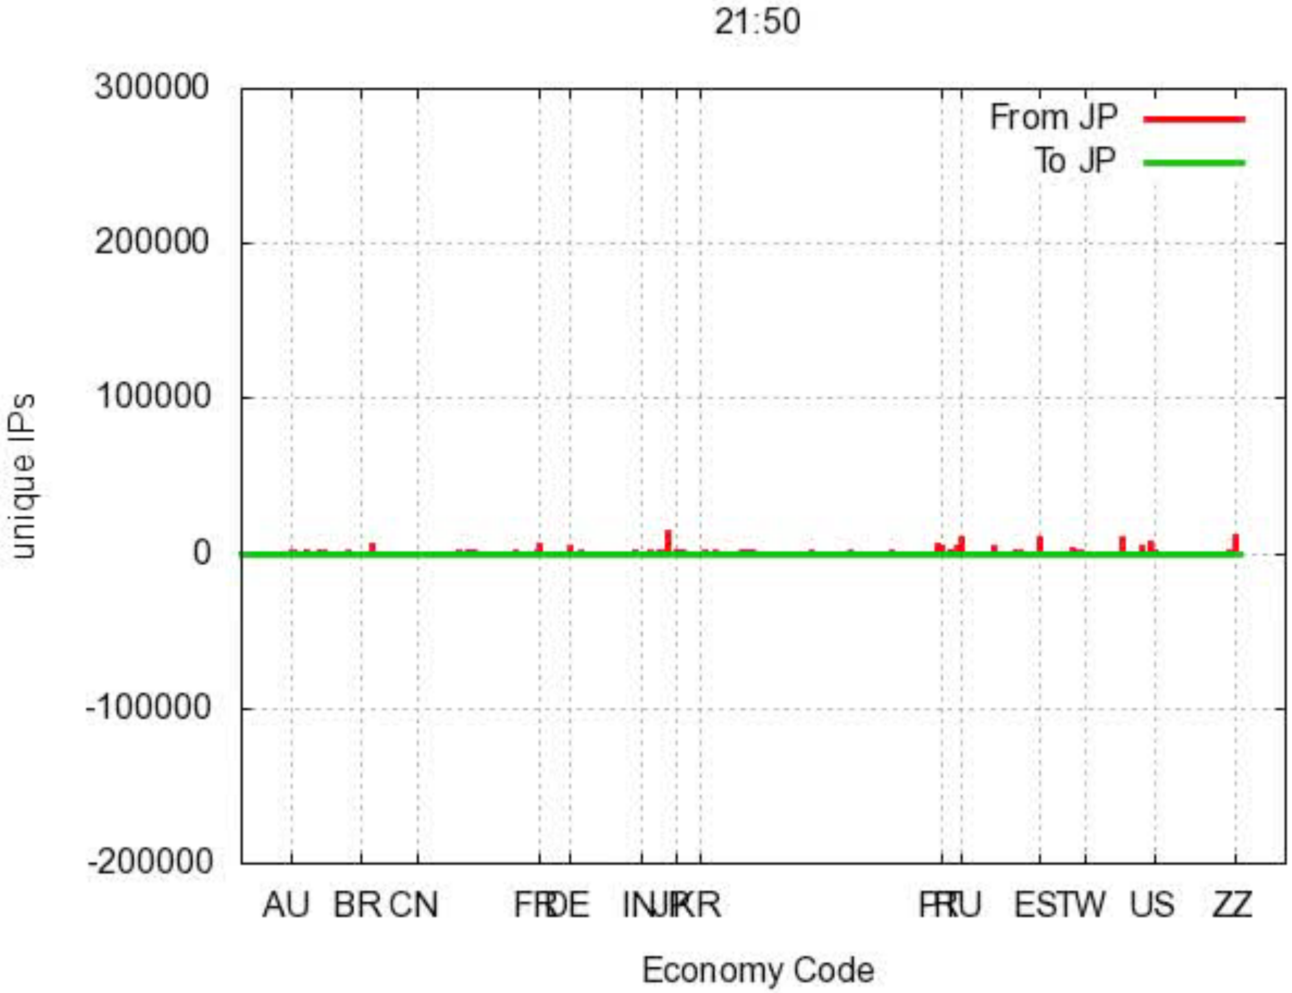

Observations

- CN, TW signals show up well
- Some phasing in the in, out {JP,CN} and {JP,TW} visible
- JP a net querier more than queried about (if you exclude US sources)

Everyone looks at Russia

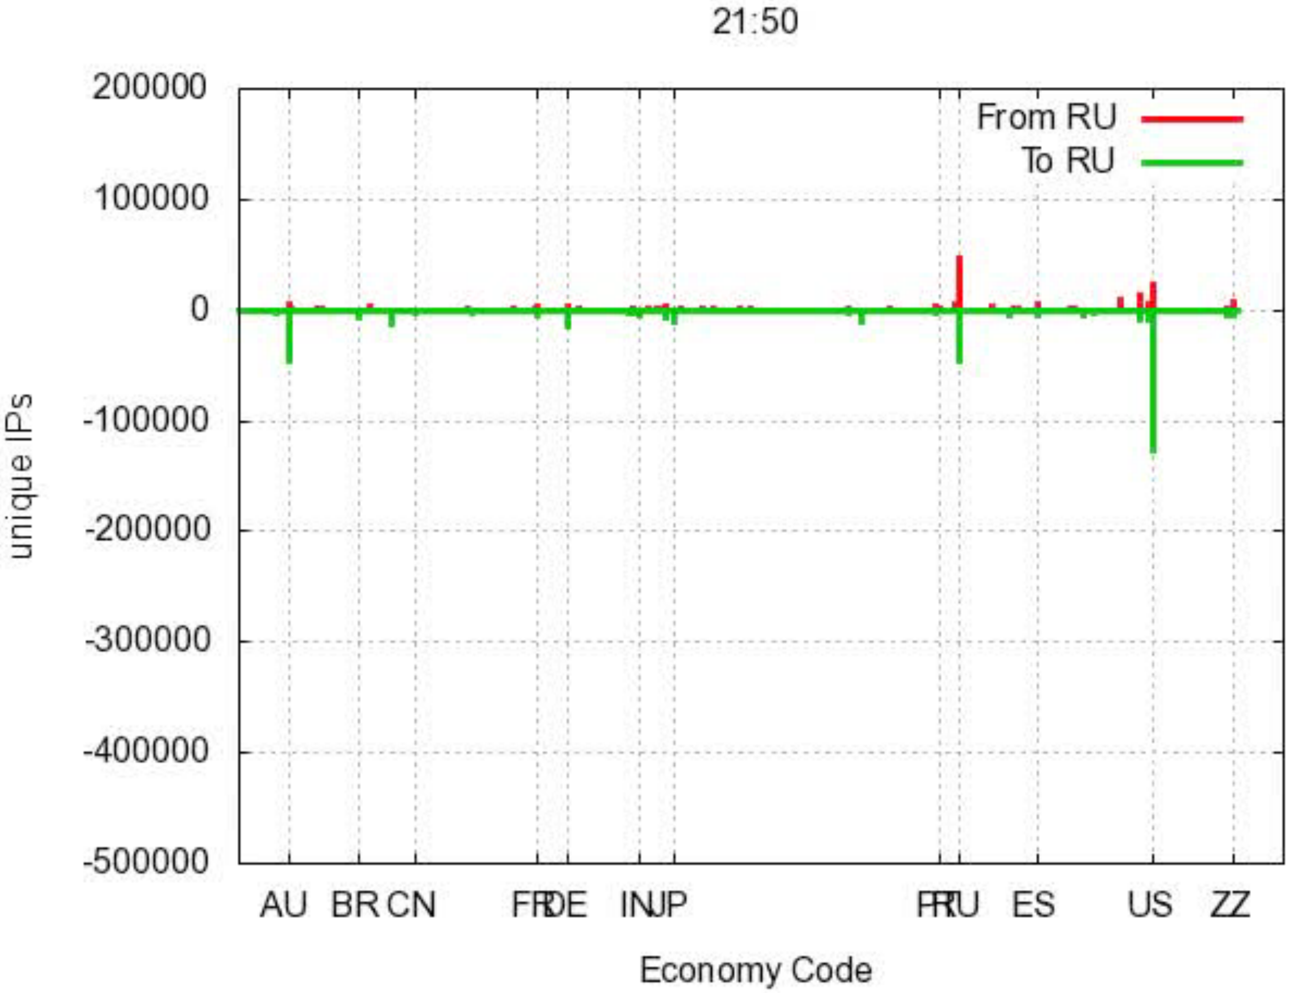

China looks at Itself

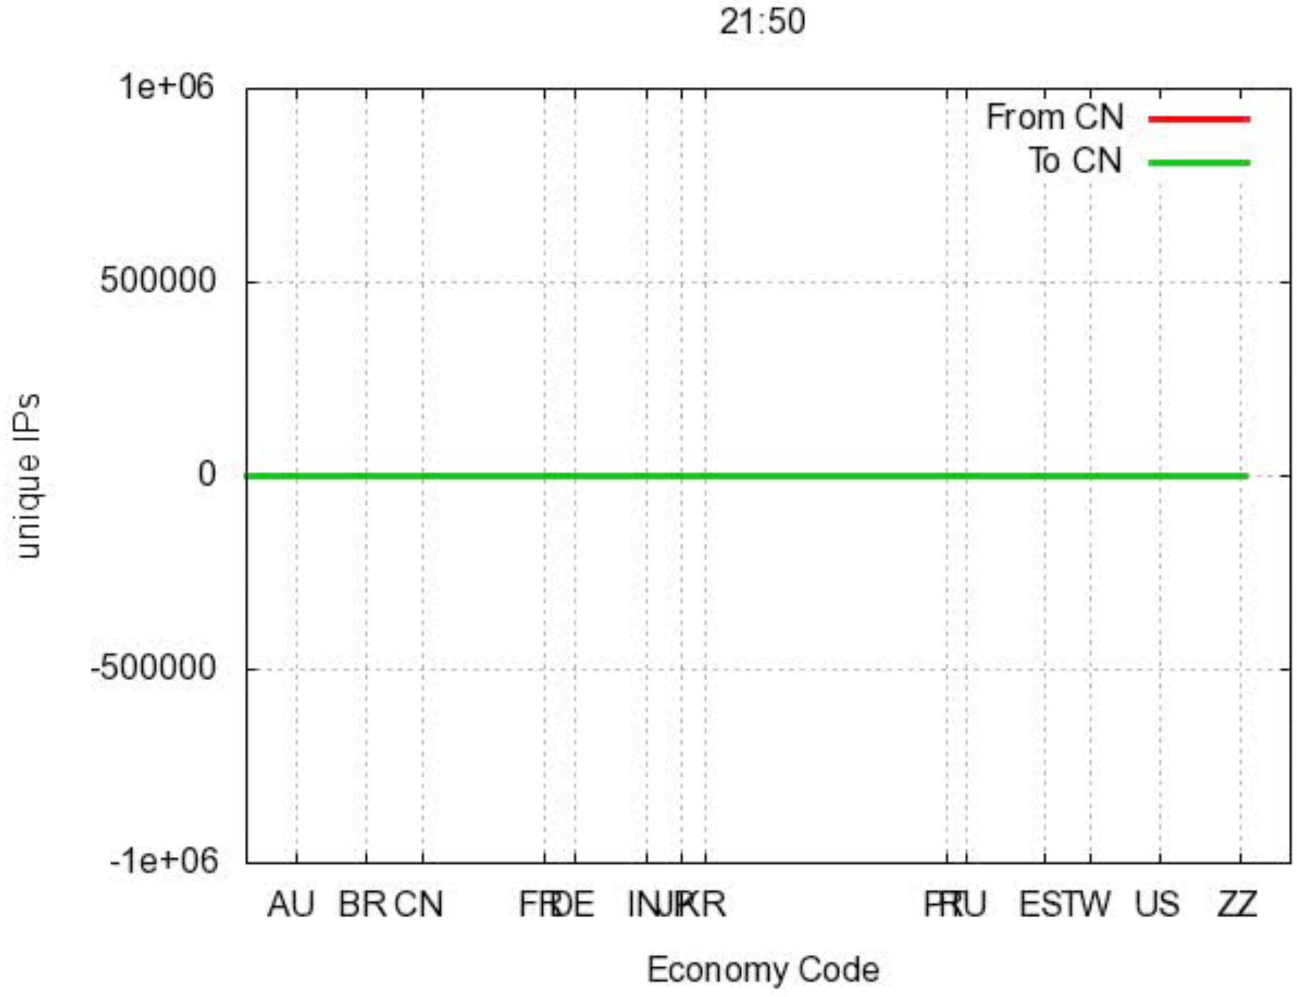

Observations

- Strong {CN,CN} signal
- More inflows than outflows

Everyone looks at India

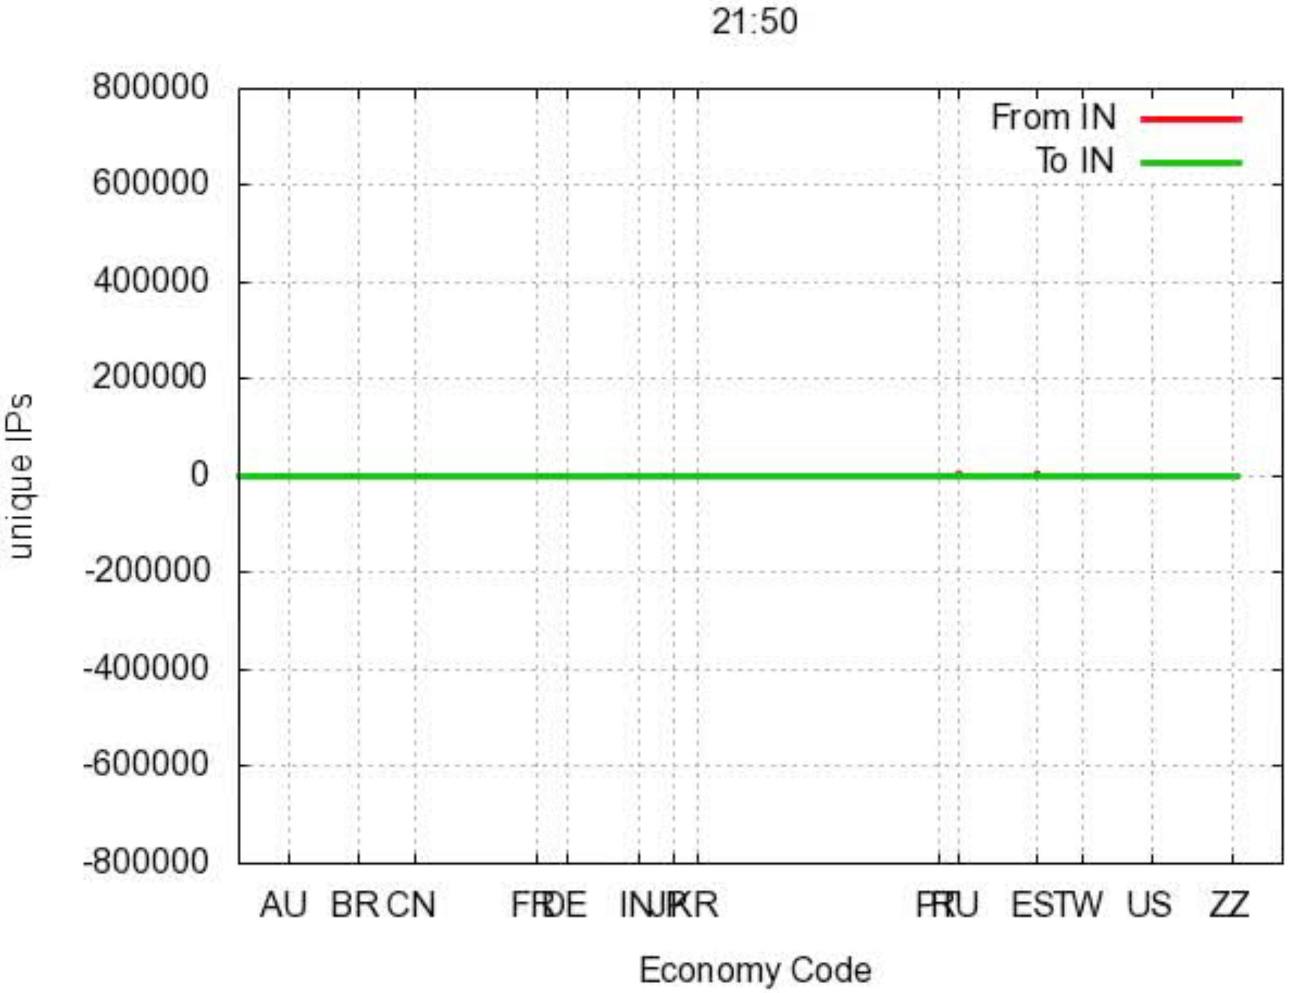

Observations

- Very little DNS flows from src=IN
- Not much interest in reverse-DNS in IN ISPs?
- Large inflows from range of economies
- (see Germany)
- Is this ‘outsource’ phenomena?

Germany looks at India

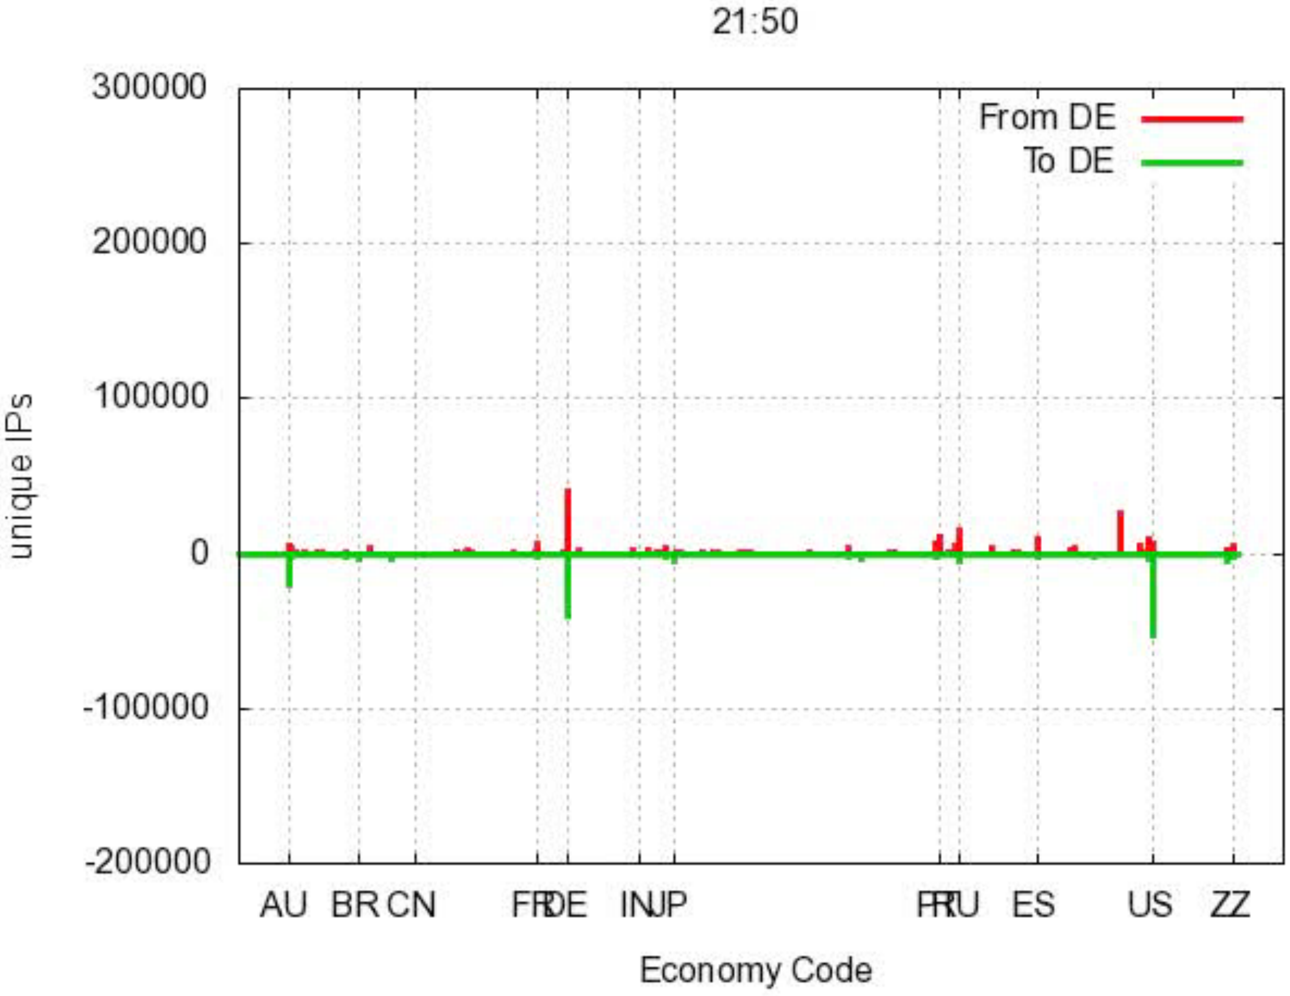

DITL 2010 goals

- Same methodology
  - Less mistakes (hopefully)
- Increased footprint of AP reverse NS
  - RIPE, ARIN may also be available (secondaries of AP in-addr/ip6 .arpa)
- DNSSEC deployment in reverse-DNS likely
  - Permits before/after analysis
- Depending on timeline, may have other NS
  - Changes in delegation tree

What is DITL/APNIC telling us?

- Reverse-DNS has interesting dynamics
- Reverse-DNS delegation is reasonably concentrated
  - Small number of servers, covering large address ranges worldwide
- Therefore (arguably) reverse-DNS is an interesting viewpoint into worldwide trends in DNS
  - Evidence of emerging inter-economy differences, similarities. Cuts across economic/developmental views
- If we can only correlate this into the real world!

Extra slides

DITL 2008-2009 AP region

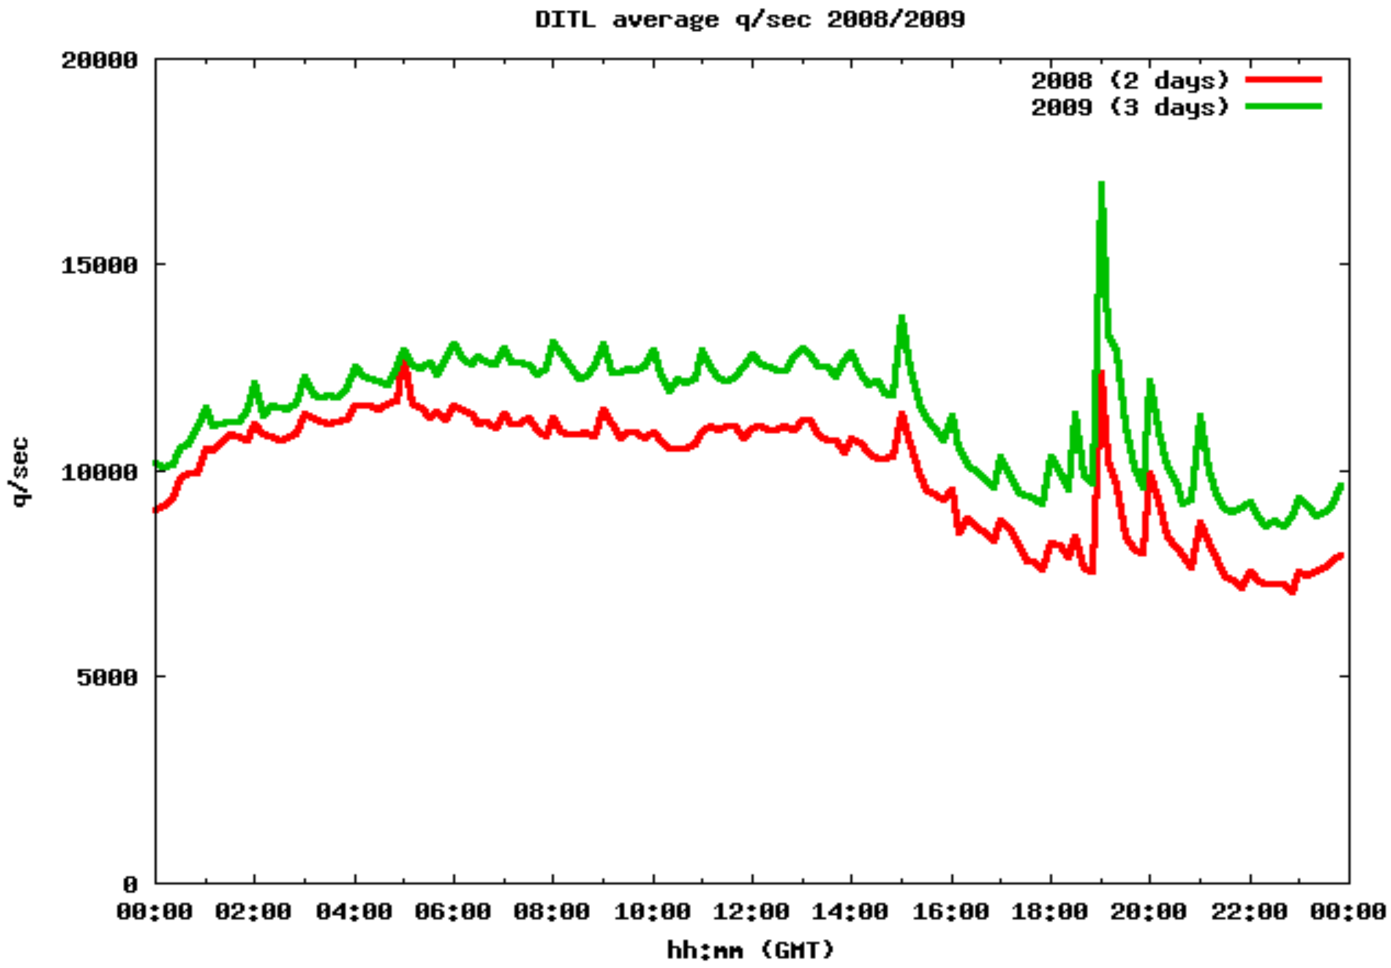

Observations

- Regular time-synced peaks
  - Cron or other periodic processing
- That ‘mega’ spike is consistent behavior 2008-2009
  - Whatever it is, its been happening for a long time

E. Asia in the World of DNS

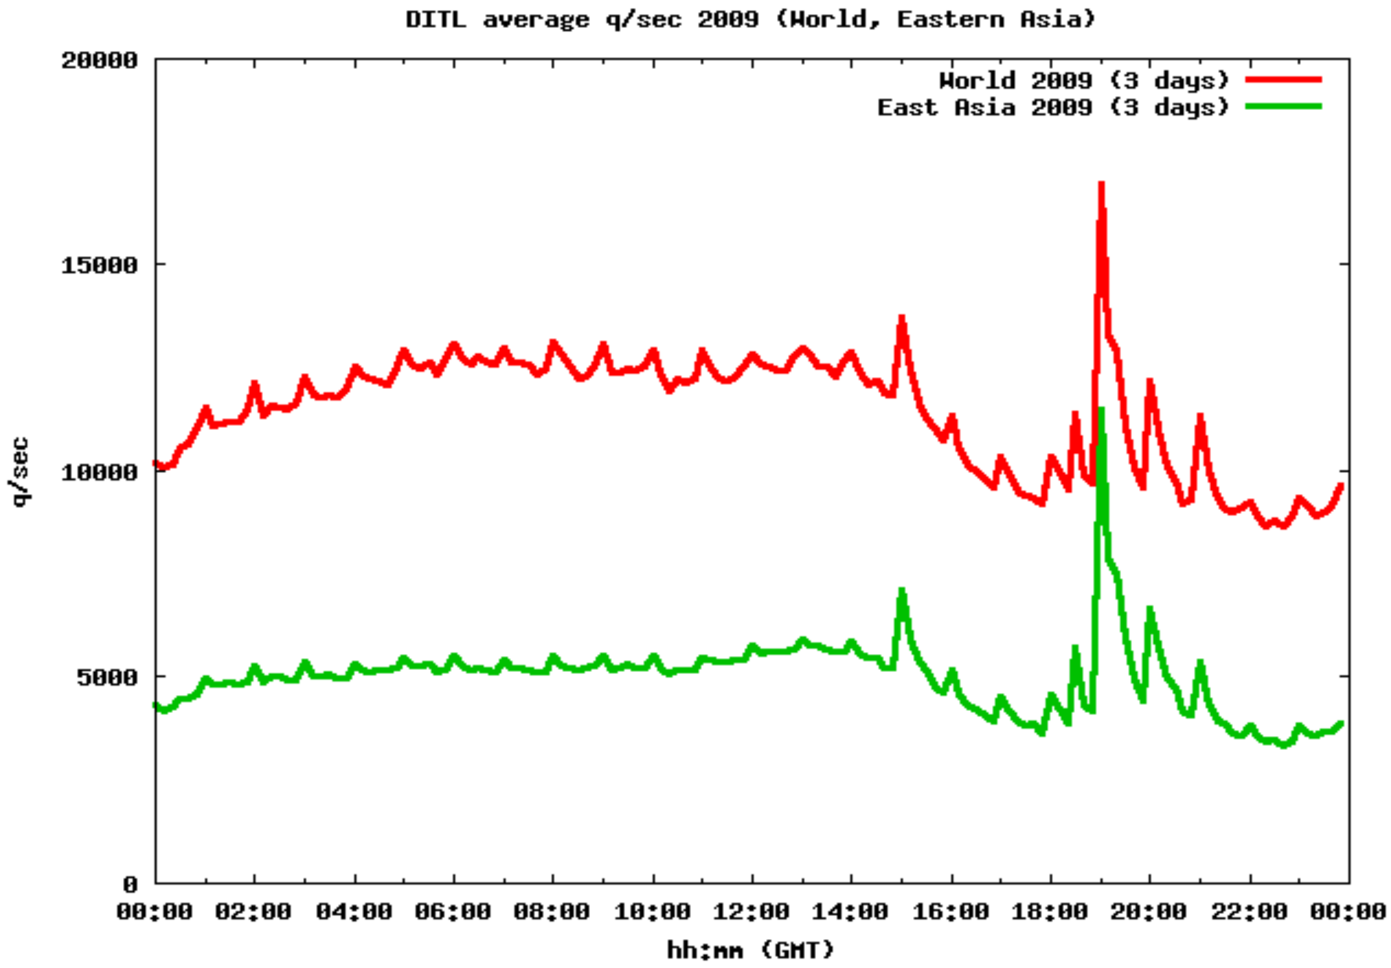

E. Asia breakdowns

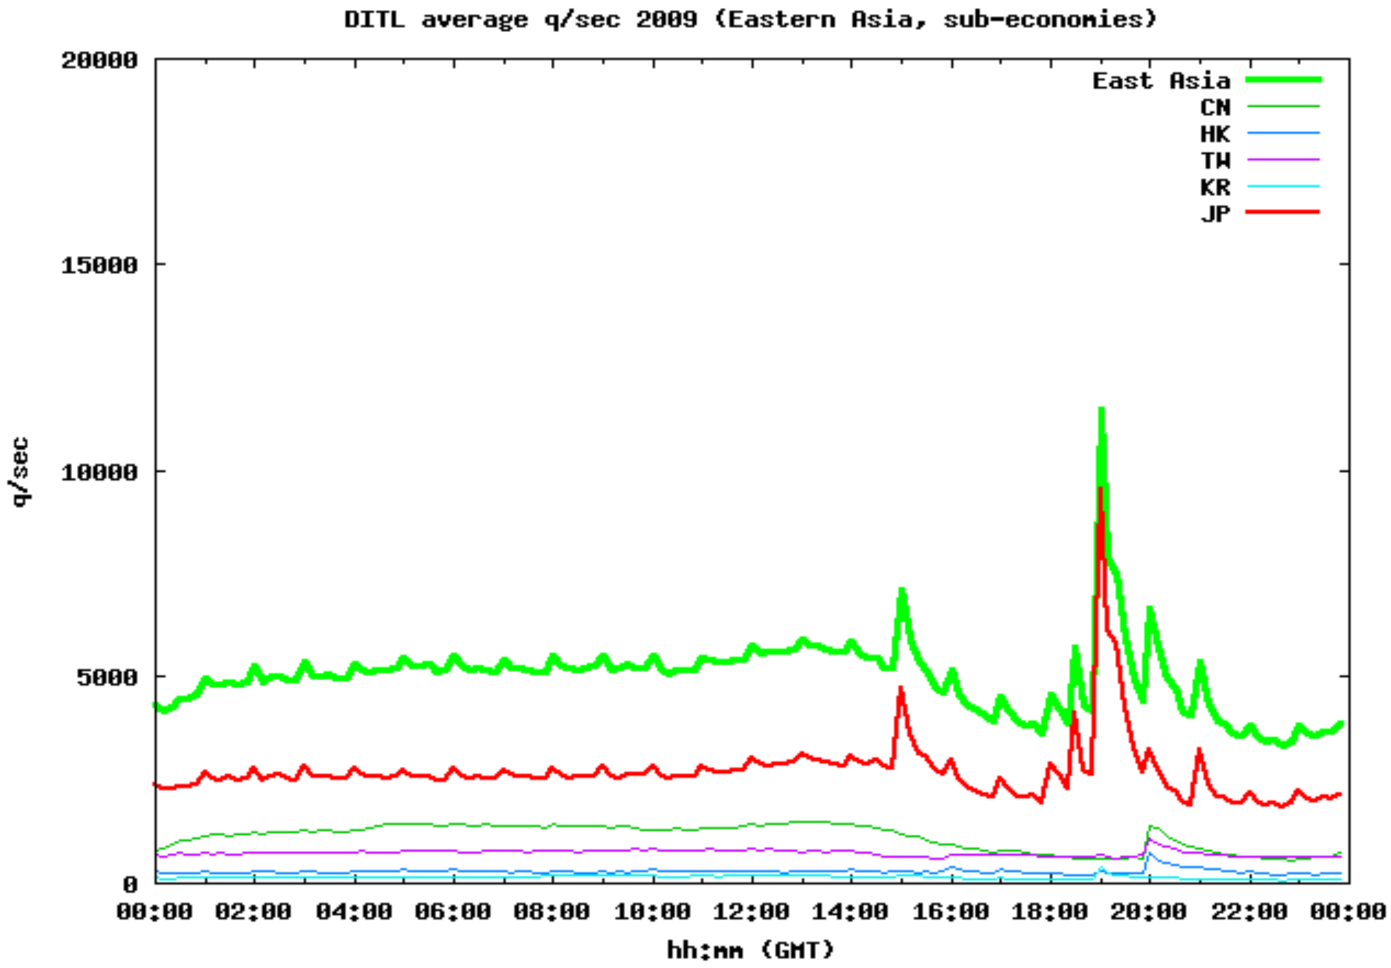

Observations

- Its Japan
- But in the ‘decomposed’ view you can also see other Economies do their own peak-fetch cycle
- HK/KR

Asia & Rest of World
Time shift

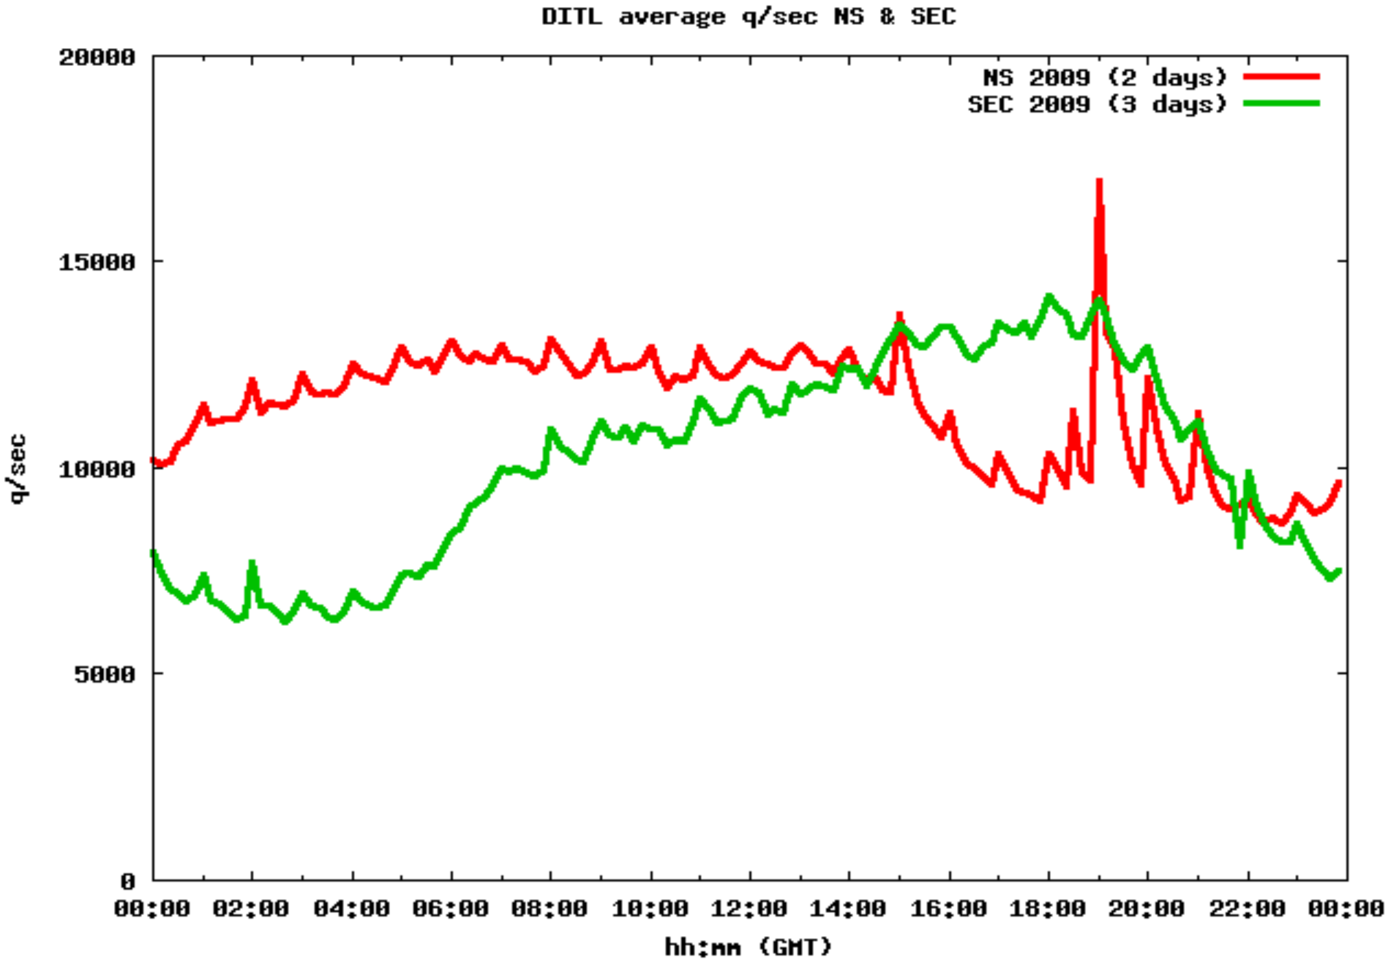

Observation

- Strong regional ‘message’ in the phasing
- Time-sync peaks seem to line up nicely

Day Samples

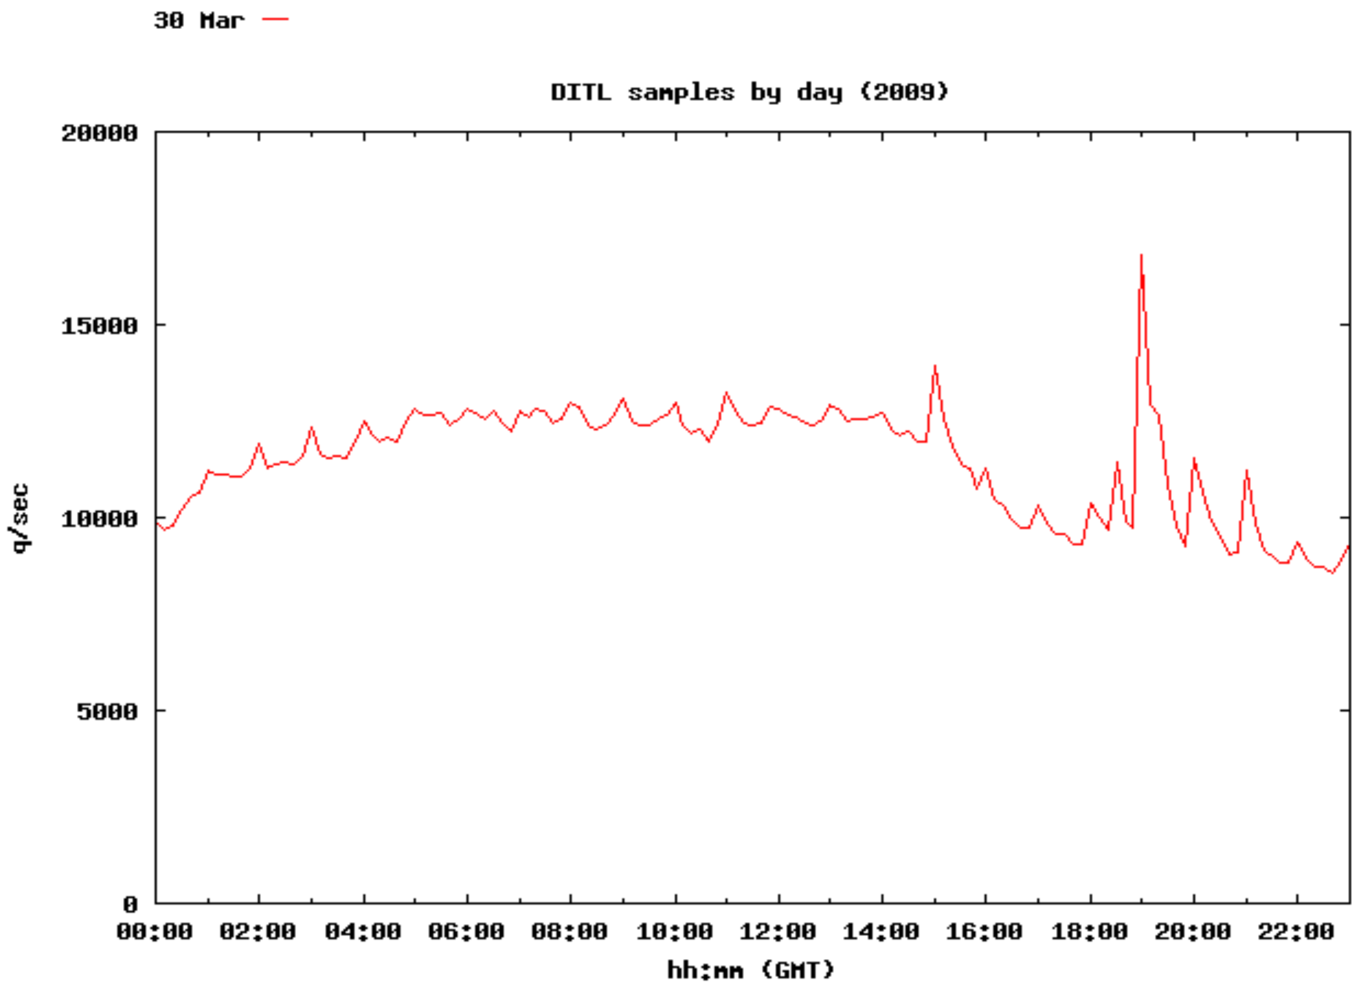
30 Mar
DITL samples by day (2009)
20000
15000
10000
5000
0
q/sec
00:00
02:00
04:00
06:00
08:00
10:00
12:00
14:00
16:00
18:00
20:00
22:00
hh:mm (GMT)

Day Samples line up

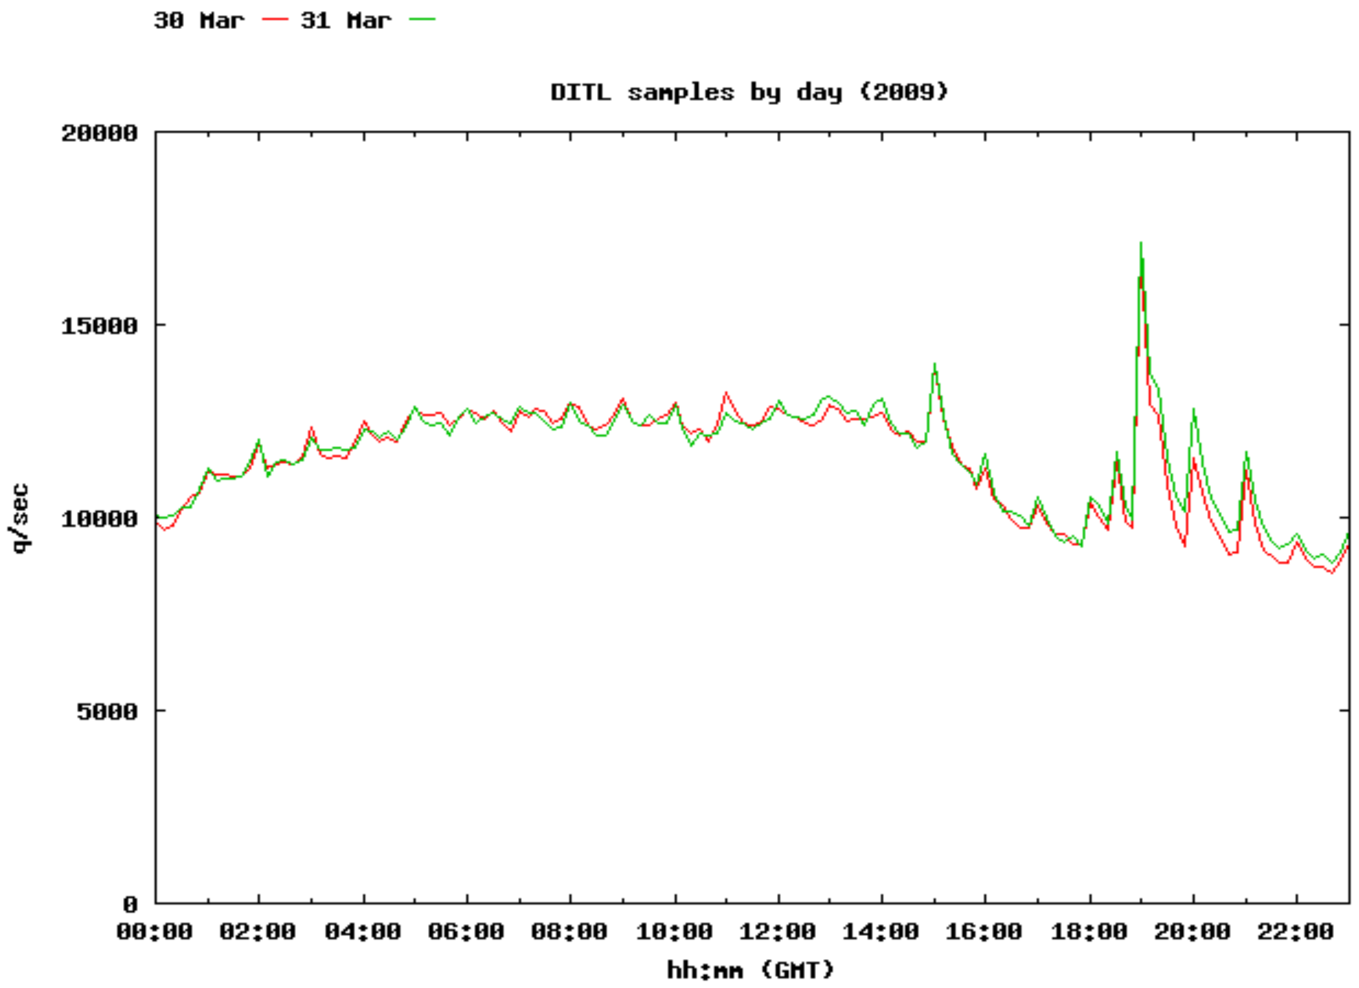

Day Samples line up strongly!

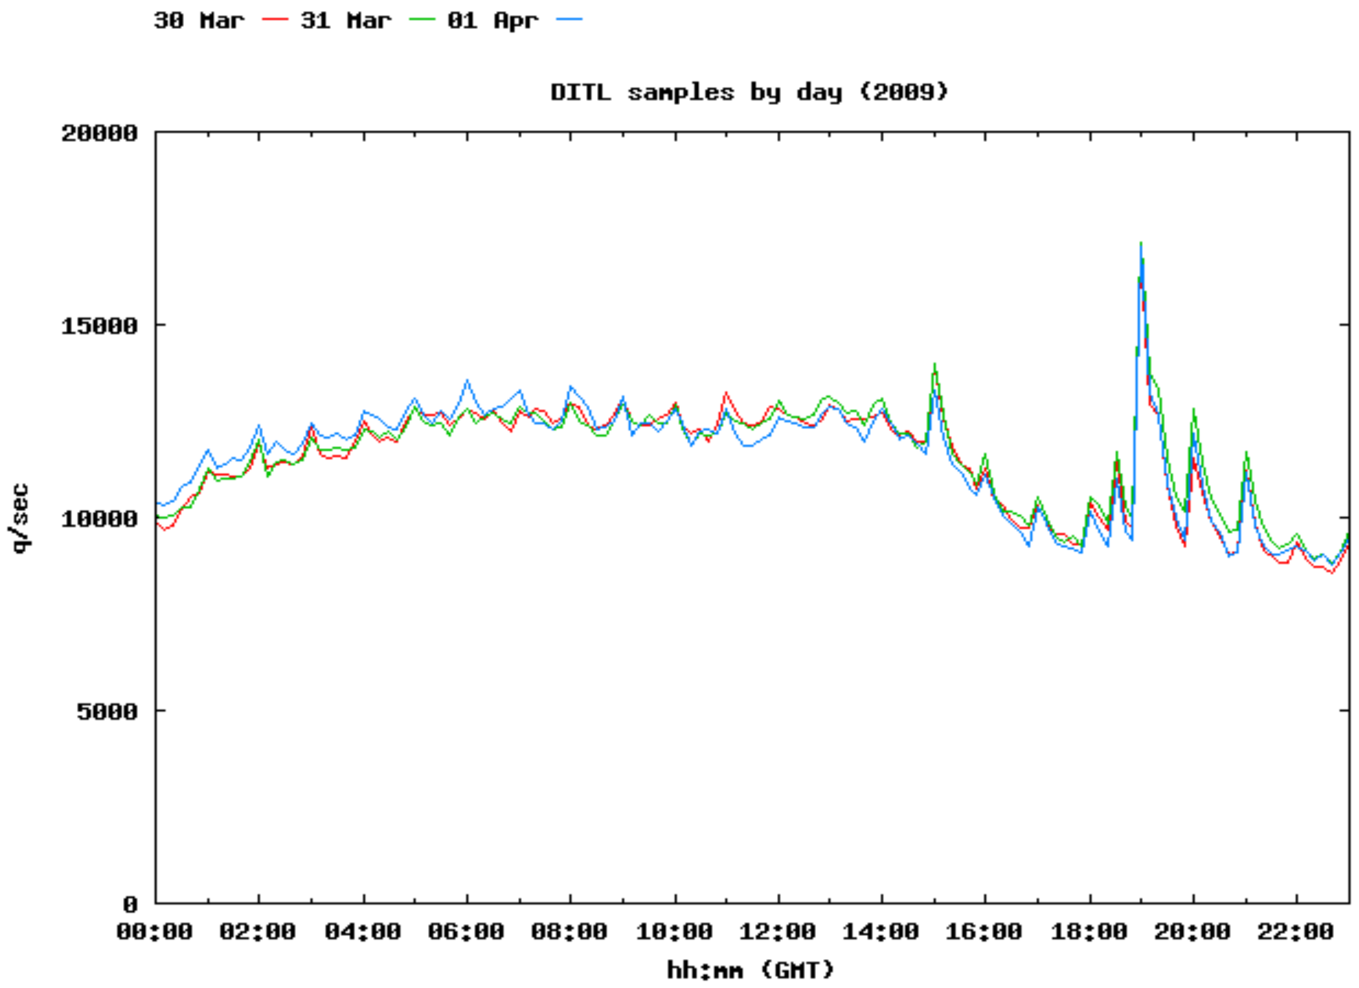

Average shows core 'shape'

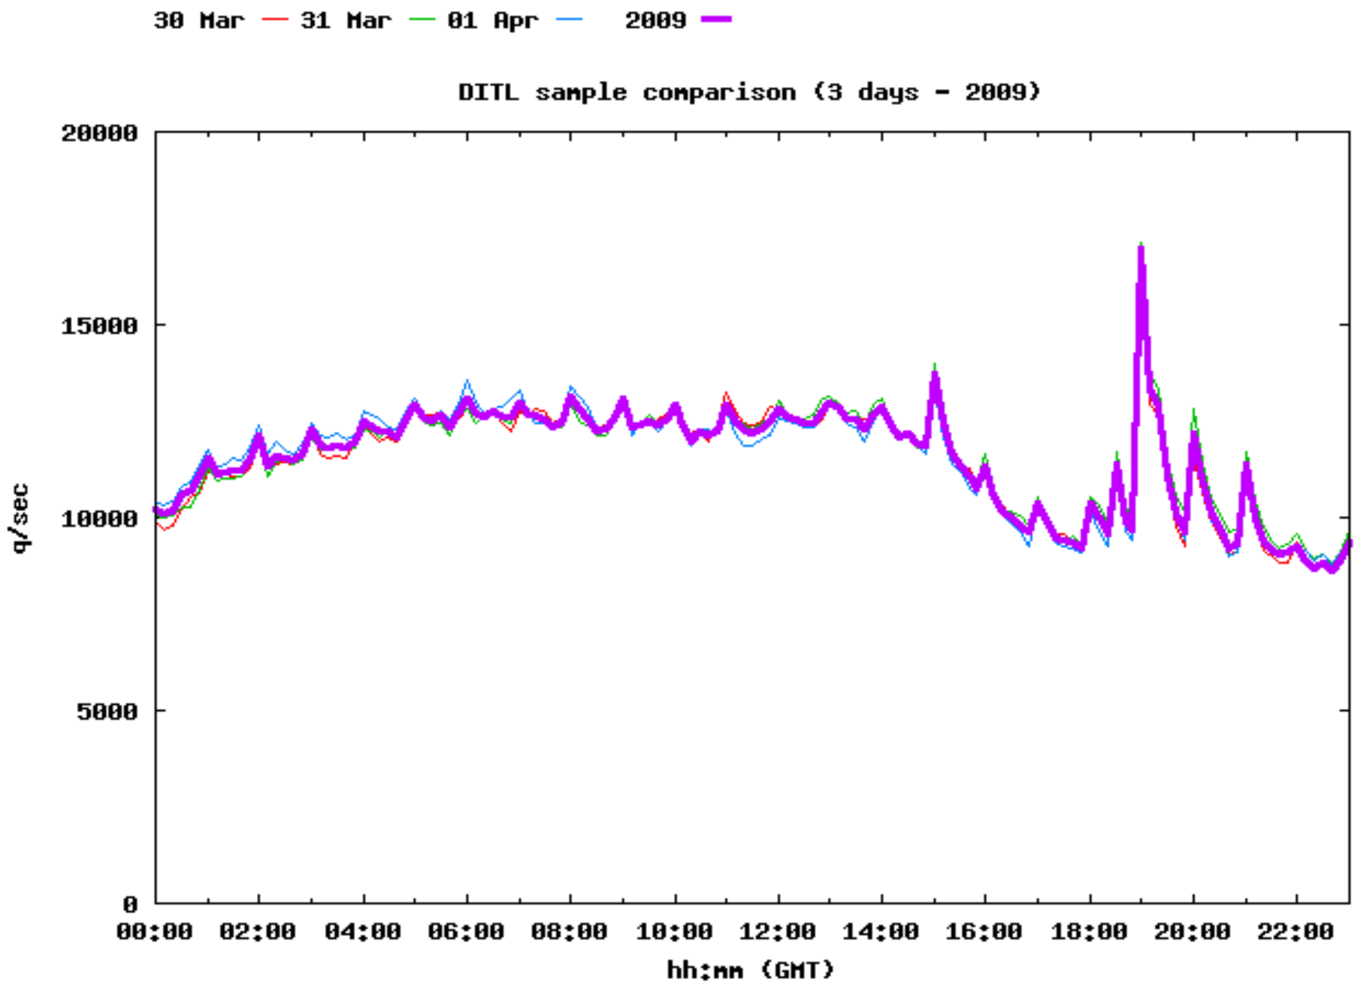

Result: year-on-year trends

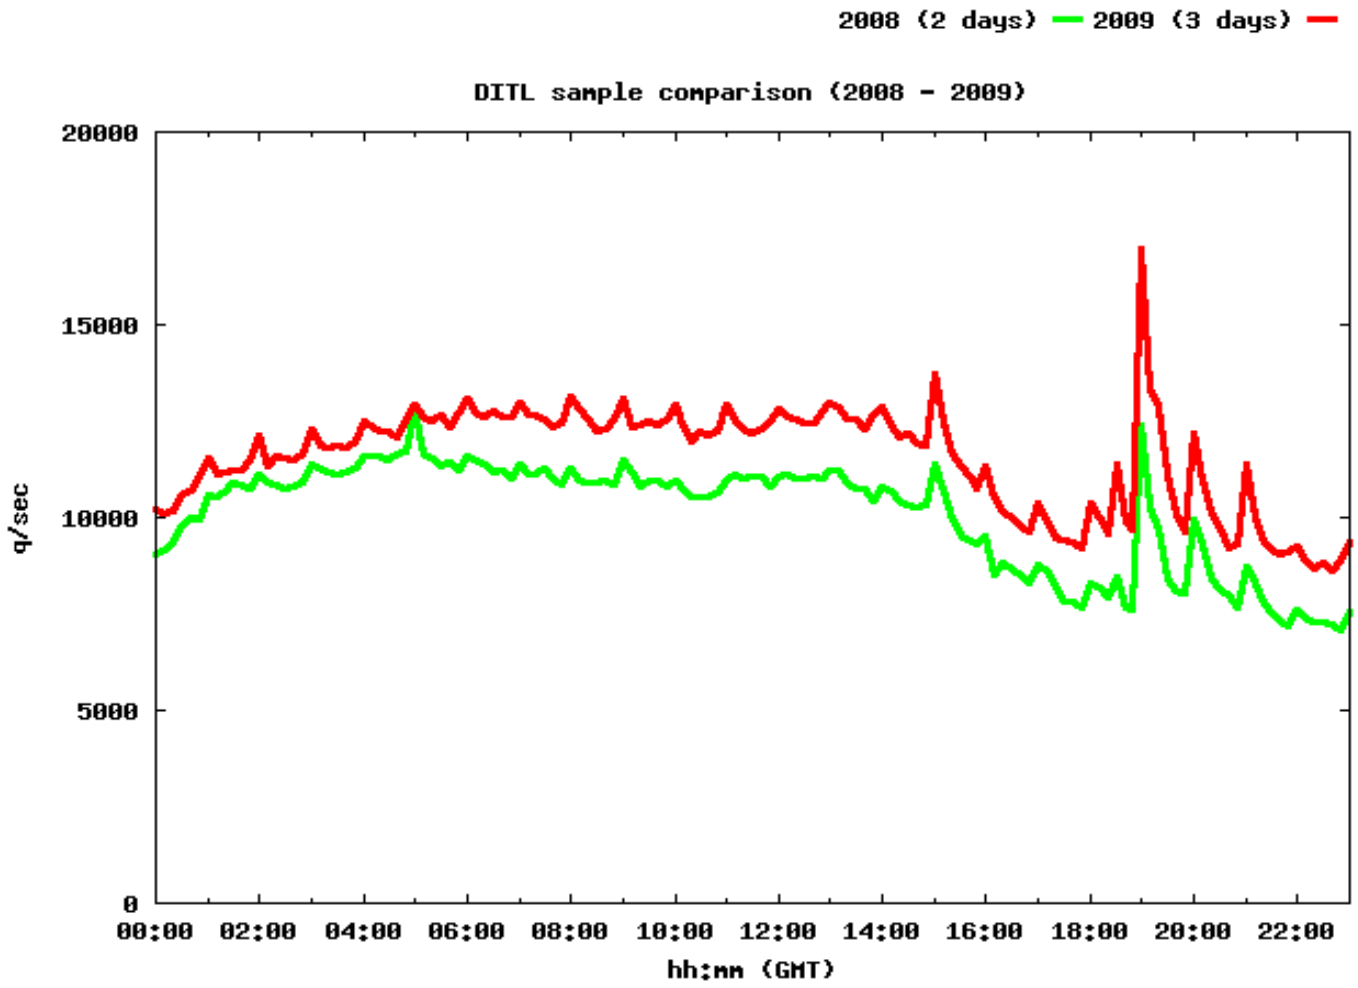
2008 (2 days)
2009 (3 days)
DITL sample comparison (2008 - 2009)
q/sec
00:00
02:00
04:00
06:00
08:00
10:00
12:00
14:00
16:00
18:00
20:00
22:00
hh:mm (GMT)