SDE Seminar: Network Measurement

U Auckland's Campus Network

and the

global Domain Name System

Computer Science, Auckland, 13 Aug 2008

Nevil Brownlee

Overview

- I’m giving this seminar on behalf of SDE
  - it’s intended to be *representative* of work in our group
  - but of course, I’ll only look at projects I’ve been directly involved with
- I’ll talk about:
  - overview of our Campus Network’s topology
  - performance of UA web servers via various ISPs
  - response times for top-level nameservers
  - the ‘DNS History Database’ project

1. U Auckland campus network (schematic)

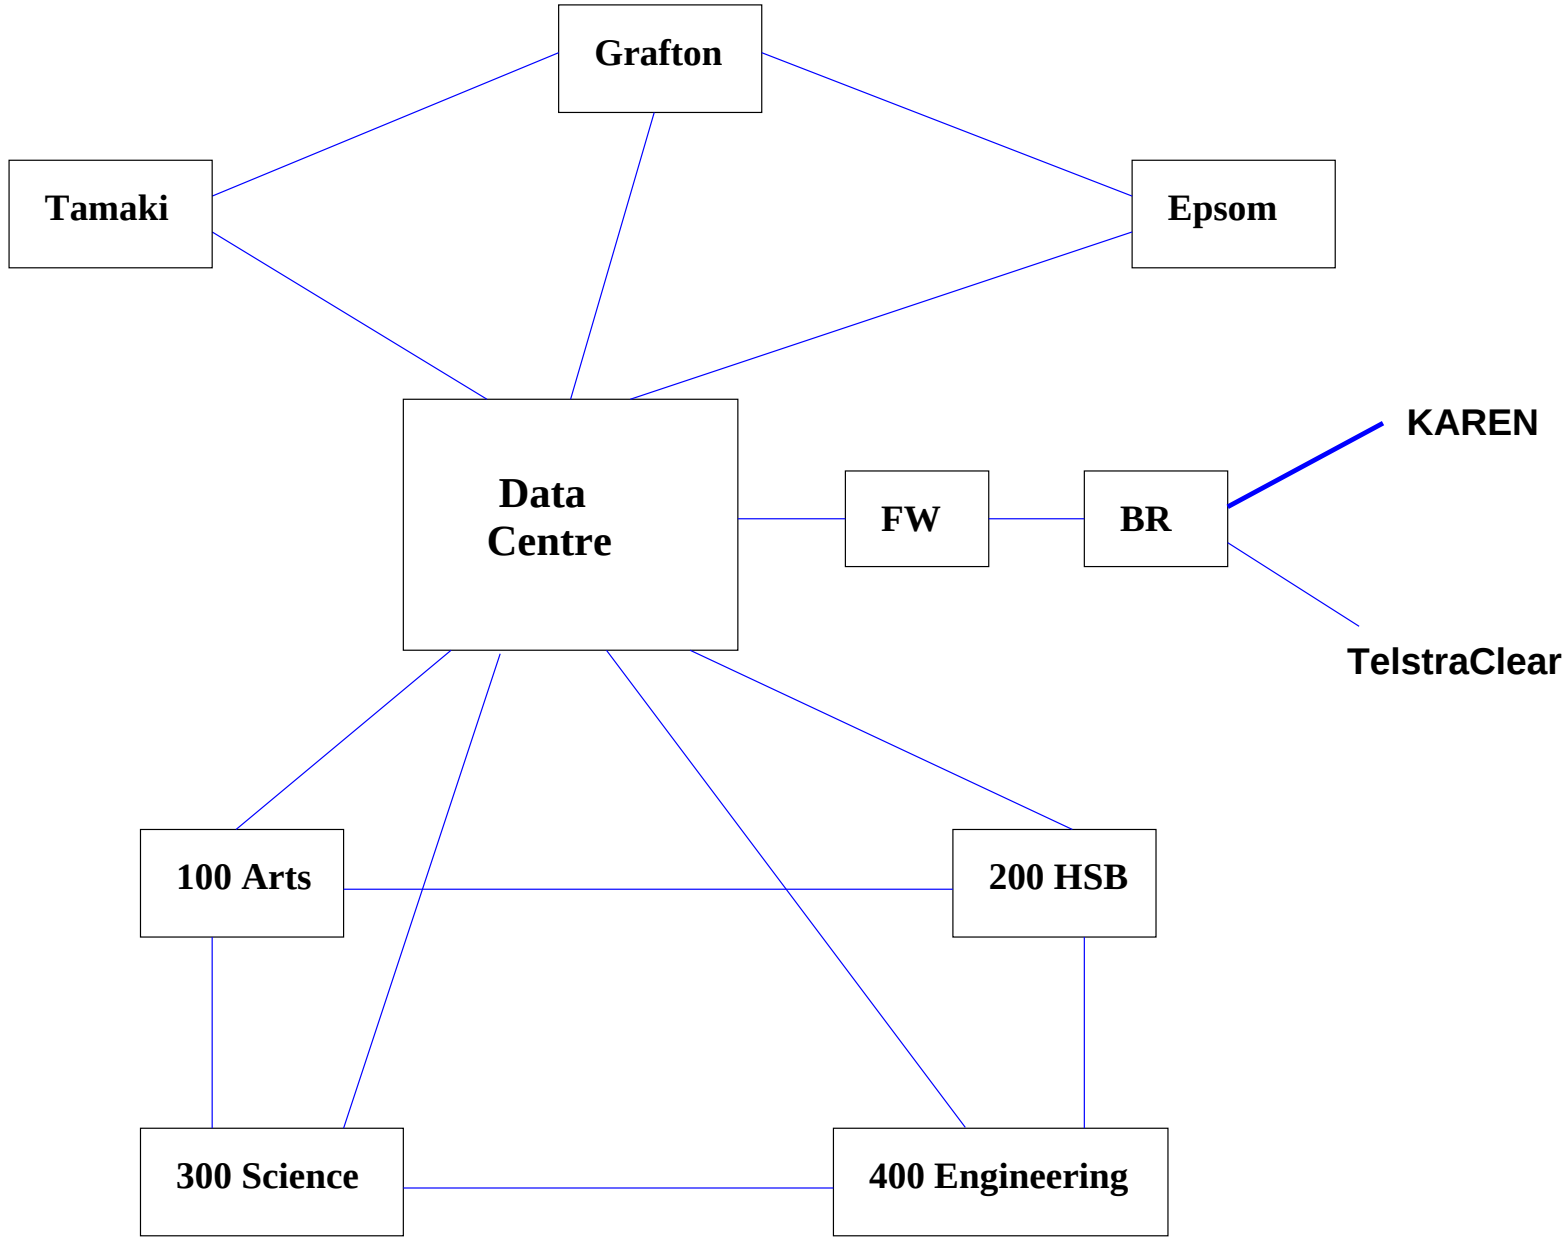

Measurement Hosts

- **nevil-res1.itss** and **nevil-res2.itss**
- Two 1U PCs, each with
  - 1TB disk, four Ethernet ports
  - one DAG card
- Located in OGG Data Centre, between our firewall and our border router
- Can observe packets going to/from Internet
  - through the border router
  - via KAREN or TelstraClear
- Used by Nevil and SDE postgrad students

Meters to observe U Auckland Internet traffic

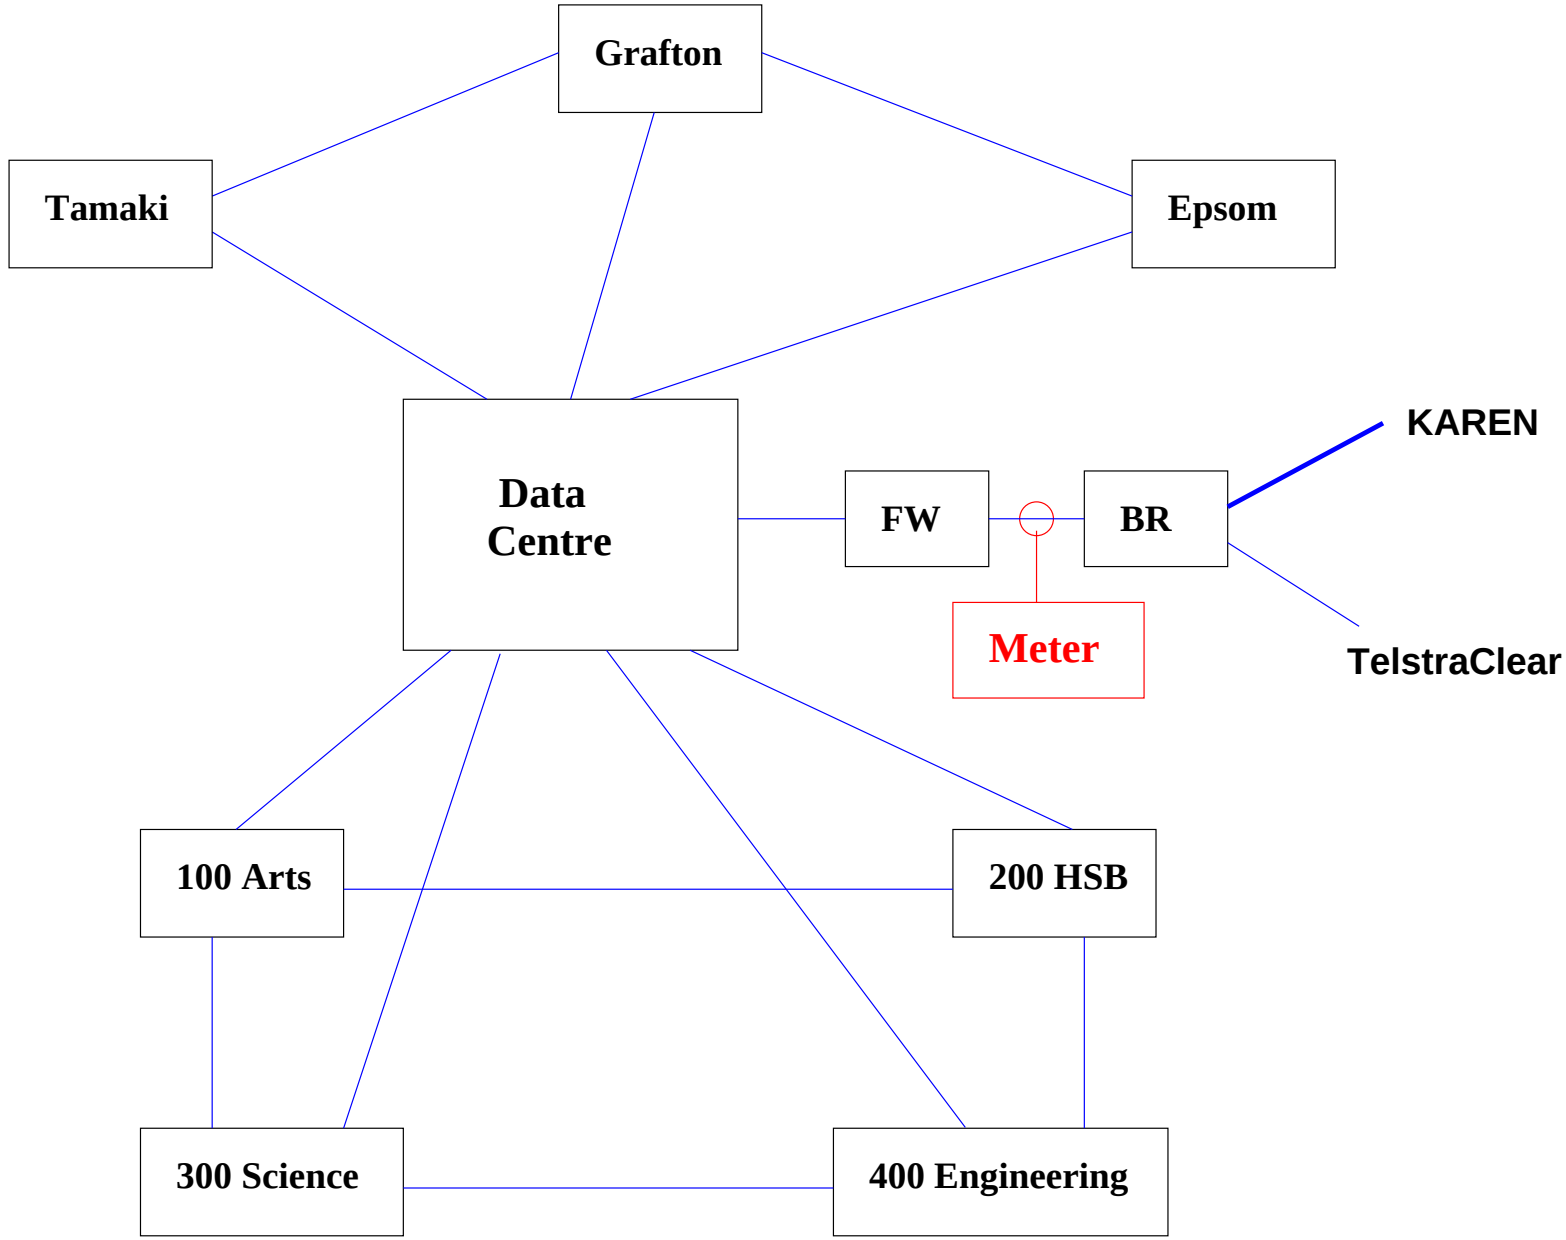

KAREN: NZ's Research and Education network

- Run by REANNZ
  - REANNZ members are Universities, CRIs, National Library
  - starting to connect schools
- 10 GB/s backbone in New Zealand
- 622 Mb/s link to U.S., 155 Mb/s link to Australia
- http://weathermap.karen.net.nz/

2. ISP Performance Monitoring

- Project initiated (and funded) by ITS
- Goals:
  - find out what proportion of users have ‘high-speed’ Internet access, i.e. better than dial-up
  - measure performance of U Auckland web servers
    - As seen by users of various ISPs
    - Using only passive measurements

Li Li's MSc project, 2005

- Goal: determine "what would suitable as ISP Performance metrics?"
- Infrastructure work
  - need to build an ISP address table,
    - Start by assuming we know /24 address prefix
    - Aggregate up as we get adjoining prefixes. [Future work: implement dynamic ptrie search for longest match]
  - connection type: couldn't measure, guess from FQDN
  - metrics: *download rate, RTT, loss rate*
    - *Connection RTT:* can't assume 3-way handshake succeeds first time, but . . .
    - Can use SYN-ACK RTT if no retransmissions
    - *Download rate:* measure over each stream, average per connection/user/ISP
    - *Loss rate:* data was too noisy for a reliable measurement

Fei Liu (Jonathan)'s Diploma project

- Nevil wrote code for **uspmon.rb,** running on nevil-res2
- Use RRD database tool to collect the data and aggregate it for day/fortnight/month/year
- http://nevil-res2.itss.auckland.ac.nz/uspmon/webperformance.php
- System has been running since Oct 06
  - download rates are rather conservative!
  - RTTs seem reasonable
  - problems with IP address changes for cecil

3. Observing the Top-level Nameservers

- Dots show server locations *in late 2001*
  - Each city lists *(root : gTLD)*, using their 1-letter names
- Geographic locations are mostly in the U.S.
- Only 13 (A..M) of each – no anycasting

Setup for Measurements of DNS traffic

- Passive data collection using *NeTraMet*
  - tap link using fibre splitter or SPAN switch port
  - write ruleset to specify flows of interest *in SRL*
  - collect flow data files
- DNS performance measurements
  - observe DNS request and response packets on link
  - meter builds distributions of request–response times
  - 5-minute stats files sent to CAIDA web site each day

Nevil's Root/gTLD Server RTT page

- Measures Request-Response times to root (and gTLD) nameservers from several locations
- Stores RTT plots in a database
- Web interface to examine data . . .
- http://www.caida.org/cgi-bin/dns_perf/main.pl

(look at week starting 19 Jul 08)

Visualising DNS RTT behaviour

- Ilze Ziedins and Ross Ihaka have been working with our DNS data
- **http://www.stat.auckland.ac.nz/∼ihaka/DNS-Times/**
- Ross has produced DNS RTT vs time contour plots ..

gTLD RTT distributions (1)

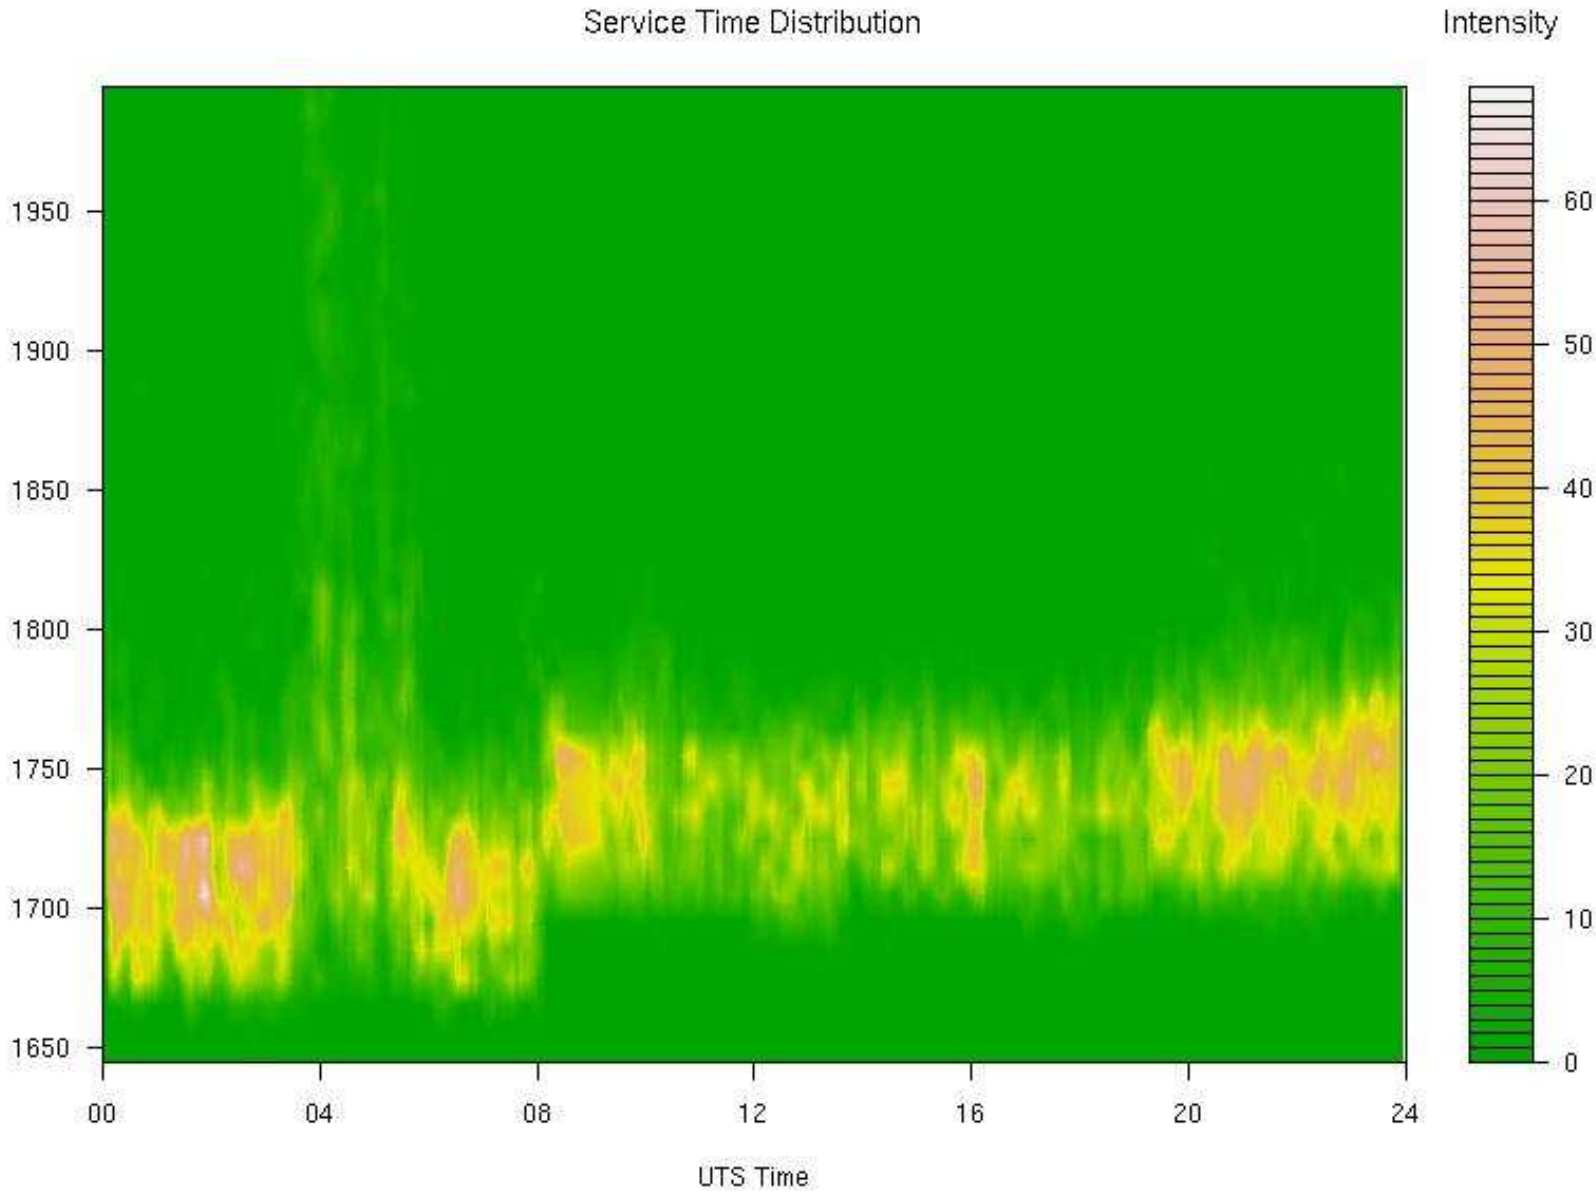

RTT/100 μs, F gTLD, observed at Auckland, Thursday 21 Aug 2003 (UTC)

gTLD RTT distributions (2)

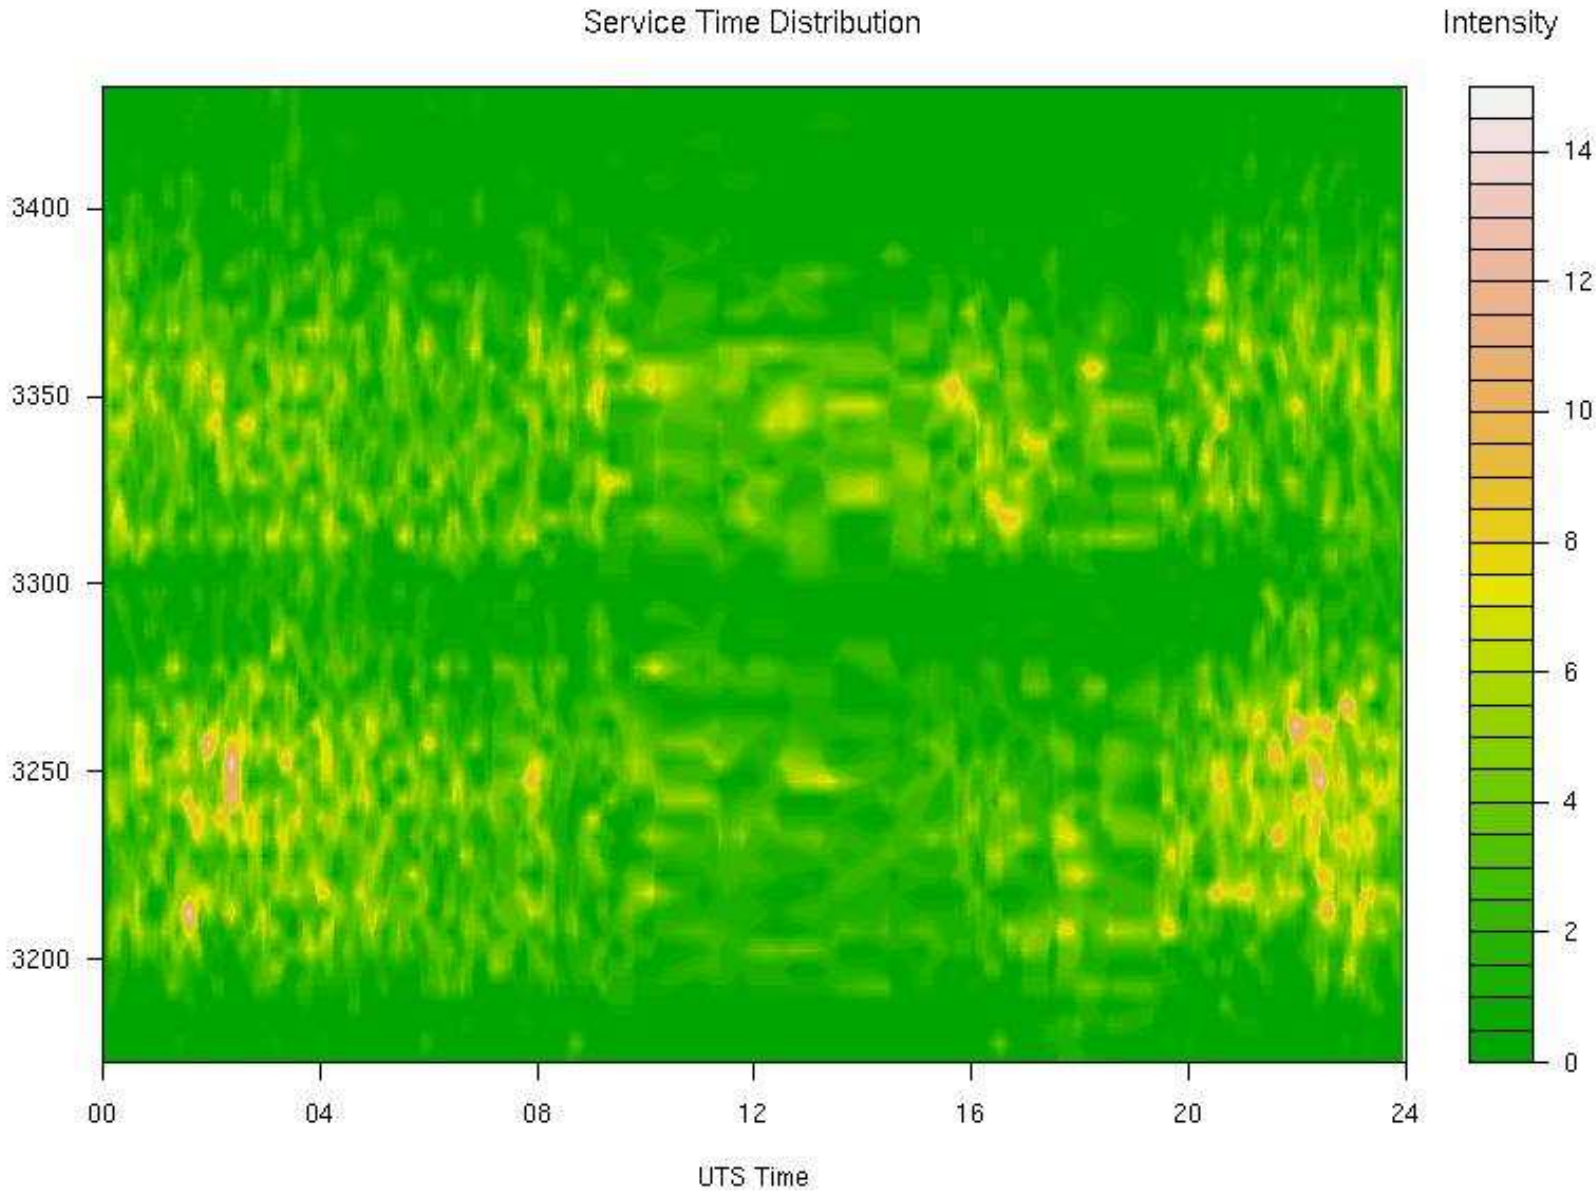

RTT/100 μs, I gTLD, observed at Auckland, Tuesday 19 Aug 2003 (UTC)

gTLD RTT distributions (3)

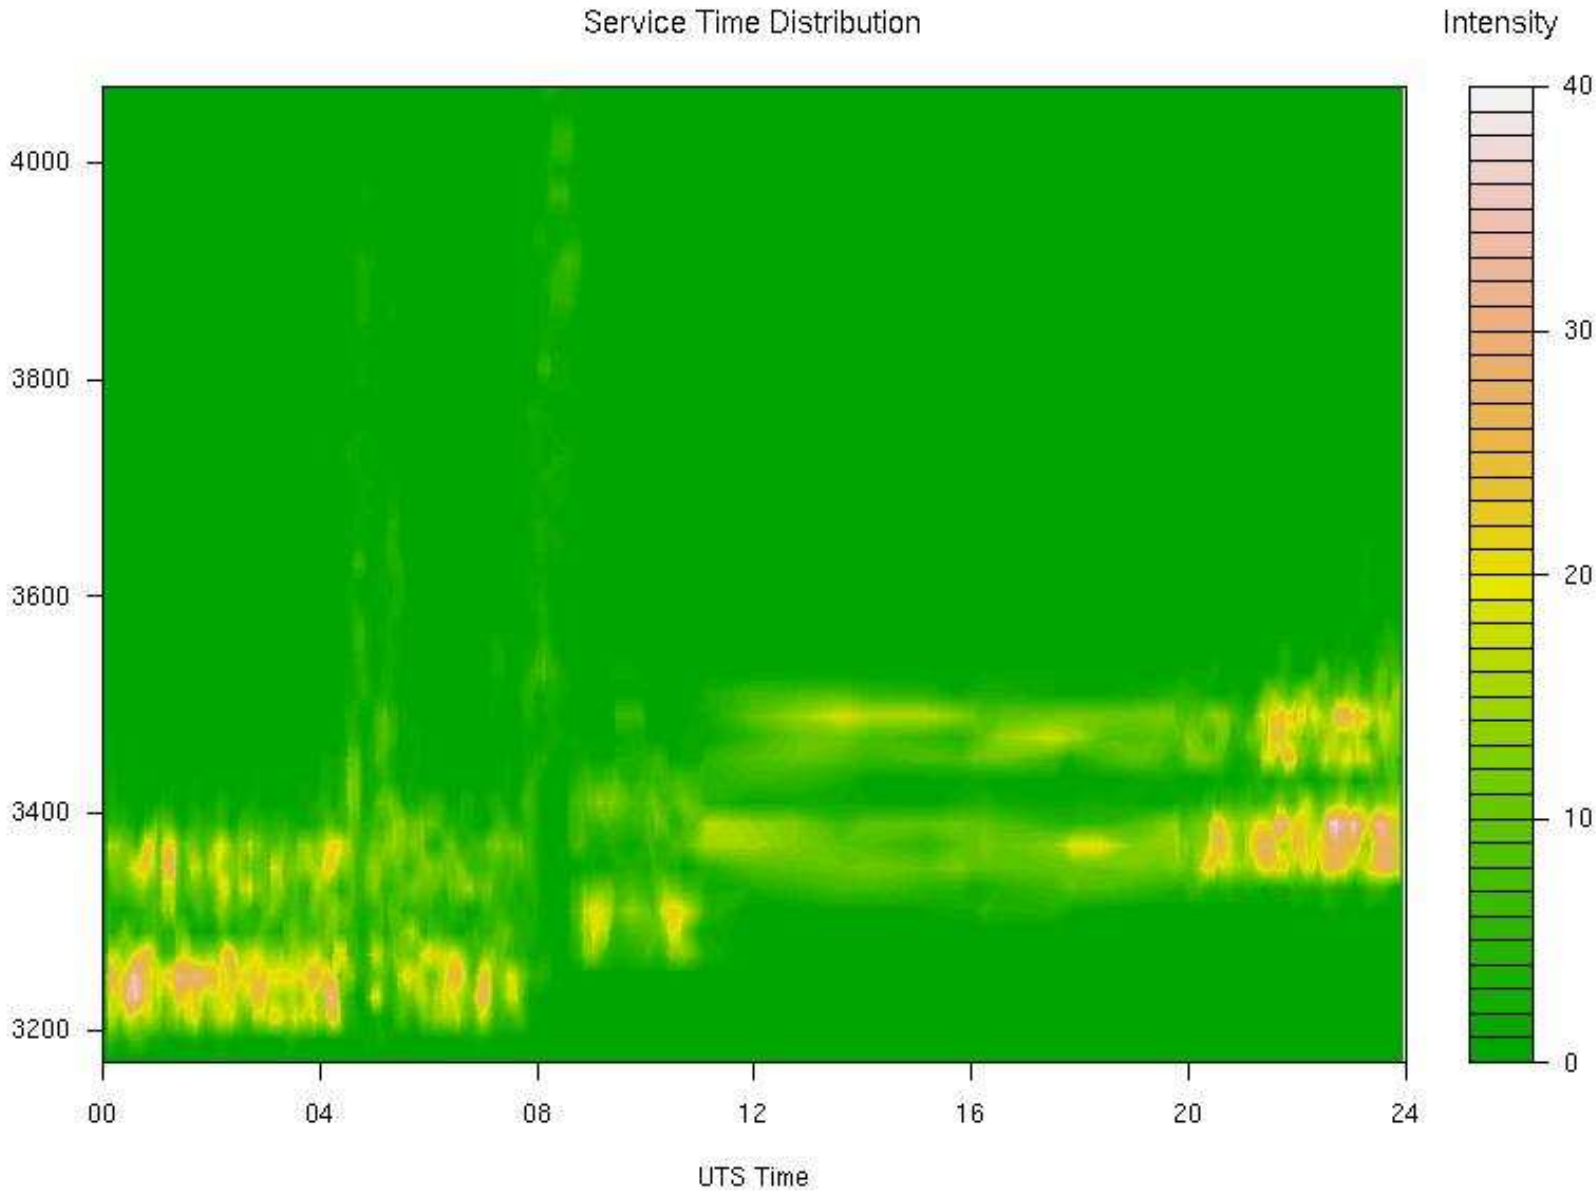

RTT/100 μs, I gTLD, observed at Auckland, Wednesday 20 Aug 2003 (UTC)

4. The DNS History Database Project, DHDB

- Started as Bojan Zdrnja's 780 project in S1 2005
- Idea is to watch DNS reply messages, and build a database of *Domain Names (FQDNs)* and the *IP addresses* they use
- Why do this?
  - attacks on DNS are becoming more common
  - for example, it's easy to capture a nameserver, and use it to divert user enquiries
  - FQDN mappings change often - when investigating network security incidents it's useful to know what address a rogue FQDN mapped to
  - DNS is distributed, each nameserver only knows about its own zone(s)
  - there's no way to retrieve the whole zone from a well-configured nameserver
  - DNS has no notion of history

DHDB: the technology

- Observe DNS packets at the edge of U Auckland's network
- Record authoritative DNS replies (the *whole* packet)
  - we only see replies to the site caching nameservers (e.g. 130.216.1.1), so there's minimal privacy risk
  - it's non-intrusive and low-impact
- At hourly intervals, process the DNS replies
  - reduce the data, push it back into the database
  - database keeps first and last time we saw an FQDN or address, along with other information from its DNS record
- Started doing this at U Auckland, plus second collector at UUNET, in Trondheim, Norway from about Nov 05
- Bojan's paper at DIMVA 2007
  - Detection of Intrusions & Malware, and Vulnerability Assessment conference

Authoritative DNS Replies at Auckland

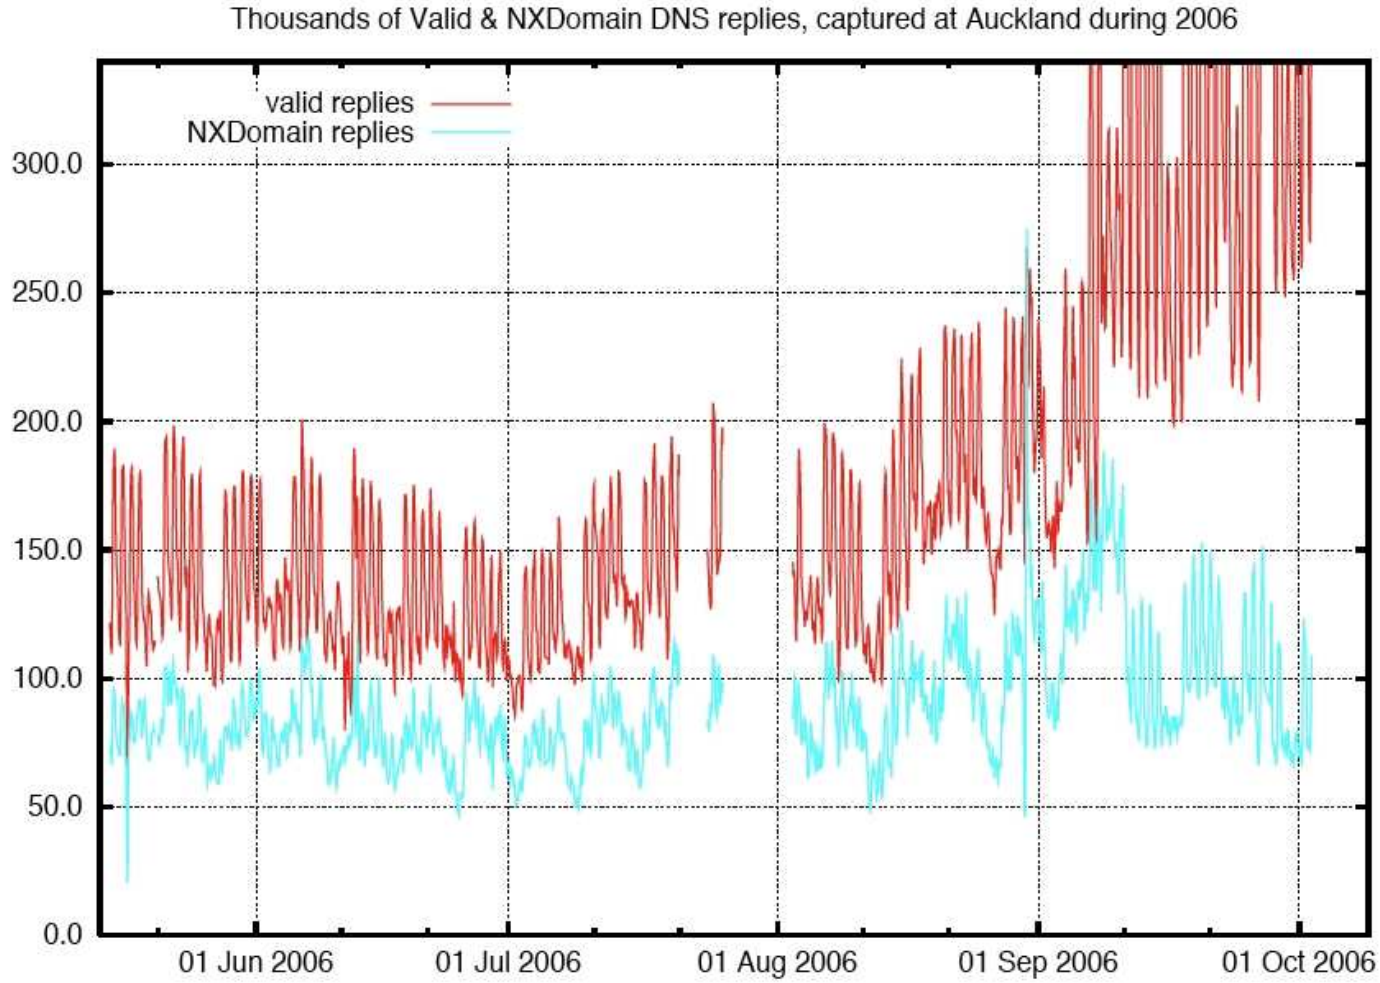

- Big rise in DNS requests in Aug–Sep 2006
- Anti-spam system was attempting to resolve all FQDNs and IP addresses it saw in email messages!

DNS Resource Records in the Database

RR type	Records	%
A	24096932	57.00
NS	757825	1.79
CNAME	652126	1.54
SOA	16281	0.04
PTR	11261024	26.64
MX	2433120	5.76
TXT	3047556	7.21
AAAA	2202	0.005
SRV	705	0.002
Total:	42267771	100

- *Our biggest A record contributors were anti-spam engines, looking up real-time black lists (RBLs)*
- *PTR records hold addresses for reverse lookups*

Typo-squatter domains

- No exploits – based on users incorrectly entering URLs
- Manual inspection revealed several big sites hosting typo-squatter web sites
- Most typo-squatting sites host hundreds of domains

DNS query	Answer RR	type
www.gmaio.com	64.20.33.131	A
openopffice.com	64.20.33.131	A
www.eikipedia.org	64.20.33.131	A
aukland.ac.nz	64.111.218.142	A
webmail.ec.aukland.ac.nz	aukland.ac.nz	CNAME

- Comment: “for ‘Auckland,’ most 1-char substitution- and adjacent-char transposition- domain names have been registered in the ac.nz domain”

Fast flux domains

- Domains with rapidly changing resource records
- Today (2006) typically used for command and control servers by bot-herders
- Characteristically have low TTL records, otherwise it takes long(er) for clients to resolve the new domain
- Easy to enumerate in the database.
  For example: **contryloansnow.com domain**

Answer	RR type	TTL	Time seen
84.105.118.33	A	5	Wed, 24 May 06 19:31:10 UTC
84.90.205.67	A	5	Wed, 24 May 06 21:11:55 UTC
86.203.193.193	A	5	Wed, 24 May 06 23:21:37 UTC

Project Status, July 2007

- Bojan Zdrnja's MSc project
- Collectors at U Auckland (New Zealand) and UUNET (Norway)
- Running on Bojan's desktop machine
- Database had about 120 million records
- Accessible at https://dnsparse.insec.auckland.ac.nz/dns
  - Username: *dimva*
  - Password: *2007*
- Strong interest from REN-ISAC in providing better hardware

Later in 2007 . . .

- Continuing strong interest
  - seven collectors
  - still running on Bojan's desktop machine
  - France CERT asked to mirror the database

- REN-ISAC were not able to actually provide assistance, alas
- Decided to develop the project into a global, long-term effort

The Developing project – 2008

- Name change to ‘DNS History Database’
  - describes what we’re actually doing
  - no confusion with earlier ‘passive DNS’ work
- Project based at U Auckland, with clear MOU
  - stable, long-term co-ordination
  - neutral party to own and oversee the MOU
- MOU sets out:
  - what project participants may do (e.g. collect data, mirror the database, analyse/summarise the data)
  - security/Confidentiality aspects to be observed

This year, so far …

- U Auckland bought a real machine to run the database
  - 1 TB of SAN disk, lots of memory
  - will gather data from the collectors, and update the database
- BeSTGRID (NZ's 'Grid Computing on KAREN' project) has also provided resources
  - VM with 2 GB of memory, and 2 TB of NFS disk
  - will run the web access server(s)
- Our DVC (Research) has agreed to sign the MOU on behalf of U Auckland
  - REN-ISAC's lawyers are deciding whether they can sign it
- Bojan has moved the database to its new machine
- Next steps
  - ask the present 'Collector' sites to sign the MOU
  - make a much better DHDB web page (on BeSTGRID site)
  - approach a few more possible Collector sites, e.g. in China

Conclusion

- I’ve presented a few network measurement projects
- Of course, the SDE group members have their own research interests and activities
- Plenty of possibilities for more seminars . . .