Infrastructure and Applications for Large-Scale DNS Data Collection

Keith Mitchell
OARC Programme Manager
Internet Systems Consortium
AusCERT
21st May 2007

Presentation Overview

- Introduction
- OARC Background
- OARC Data-gathering
- “Day in the Life of the Internet”
- Root Server Attack 6th Feb 2007

Introduction

What is the DNS ?

- Internet Domain Name System
- Provides conversion from human/application-friendly "domain names" (e.g. www.isc.org)..
- ..to network-friendly Internet Protocol addresses (e.g. 204.152.190.196, fe80::200:1aff:fe1a:2761)
- Highly distributed servers, hierarchy delegated from
  - 13 "root" servers
  - "top-level" (e.g. ".au", ".org") servers
  - providers
  - users

What is ISC?

- Internet Systems Consortium, Inc.
  - Headquartered in Redwood City, California
  - 501(c)(3) Nonprofit Corporation
- Mission:
  - To develop and maintain production quality Open Source software, such as BIND and DHCP
  - Enhance the stability of the global DNS through reliable F-root nameserver operations and ongoing operation of OARC
  - Further protocol development efforts, particularly in the areas of DNS evolution and facilitating the transition to IPv6.

What is OARC ?

- Operations, Analysis and Research Center for the Internet
- Co-ordination centre to protect Global DNS infrastructure
- Trusted, neutral environment for operators and researchers to:
  - gather and share data
  - co-ordinate response to attacks
- Secretariat run and managed by ISC

Speaker's Background

- Internet operations and development since 1986
- Network security *survivor* rather than expert...
- Founder and CTO of UK's first commercial ISP, *PIPEX* 1992-1996
- Founder and Executive Chairman of London Internet Exchange, *LINX* 1994-2000
- Founder and Director of *Nominet UK* 1996-2002
- Chair of *RIPE NCC* Executive Board 1998-2000
- Founder and CTO of pan-European commercial IXP operator, *XchangePoint* 2000-2004
- Chair of *UK Network Operators' Forum* 2005-
- Moved to US (Cleveland OH) Q2 2006-

OARC Background and Introduction

OARC Mission

- Provide trusted channels for Internet incident reporting and handling
- Facilitate confidential sharing of DNS operations data
- Interface with research community for analysis and publication
- Outreach to vendors, end-users and law enforcement

OARC Motivation

- DNS infrastructure makes everything work as expected
- DNS outage of any network service provider or large content provider affects everyone using the Internet
- Growing resource demand for Internet:
  - abuse prevention
  - infrastructure protection
  - operational co-ordination

OARC Motivation

- Increasing incidence of attacks against the DNS
- DNS increasingly implicated in and compromised by Botnet activity
- A lot of unwanted traffic on the Internet is a result of DNS misconfiguration
  - e.g. in-addr queries to RFC1918 addresses
- New DNS technology challenges
  - DNSSEC, IDN, ENUM, IPv6

DNS as Abuse Vector

- "Fast-flux" short-lifetime domains for BotNet C&Cs
  - Also for Phishing
- "Cache poisioning" injecting bogus data
- "Pharming" hi-jacking of DNS queries from infected machines
- DDoS amplification
  - unsecured recursive resolvers open to answer source-spoofed queries from anywhere
  - response packets **much** larger than queries

DNS as Abuse Victim

- Mostly large-scale attacks against top-level DNS infrastructure
  - Microsoft outage in 2001
  - DDoS attack on Root Servers 2002
  - Open recursive resolvers Q1 2006
  - DDoS attack on Root Servers Feb 2007
- Attack motivation unclear
  - proof of capability ?
- DNS infrastructure increasingly robust to these

DNS as Abuse Preventor

- By doing large-scale gathering of DNS traffic, it becomes possible to identify traffic patterns underlying abuse
- Network operators in best position to gather data
- Researchers often in better position than network operators to analyse data
- Data resource potentially available to law enforcement to trace abuse sources

OARC Members

- Current total 44, includes:
  - 6 root server operators
  - 2 gTLD operators
  - 12 ccTLD operators
  - 11 DNS implementers
  - researchers at 5+ institutions
  - RIRs, DNS registrars, operators
- 10+ potential new members in pipeline

OARC Members

- Afilias
- AFNIC
- APNIC
- Autonomica
- BFK
- Cambridge Univ
- ChangeIP.com
- CIRA
- Cisco
- Cogent
- CZ.NIC
- Damballa
- DENIC
- eNom
- EP.net
- F-root
- Georgia Tech
- Google
- II-F
- Internet Perils
- ISC
- ISoc-IL
- Microsoft
- NASA Ames
- NASK
- *NIC.CL*
- NIDA
- Nlnet Labs
- Nominet UK
- NTT
- *OpenDNS*
- PIR
- Registro.BR
- RIPE NCC
- Shinkuro
- SIDN
- Team Cymru
- UMR.edu
- NeuStar/uDNS
- UMD.edu
- WIDE

OARC Member Services

- DSC Data Gathering
  - Domain Statistics Collector
- Data Analysis
- Member-only mailing list
- Other closed DNS mailing lists
- Encrypted jabber.oarc.isc.org chat server
- https://oarc.isc.org portal

DSC Data Gathering

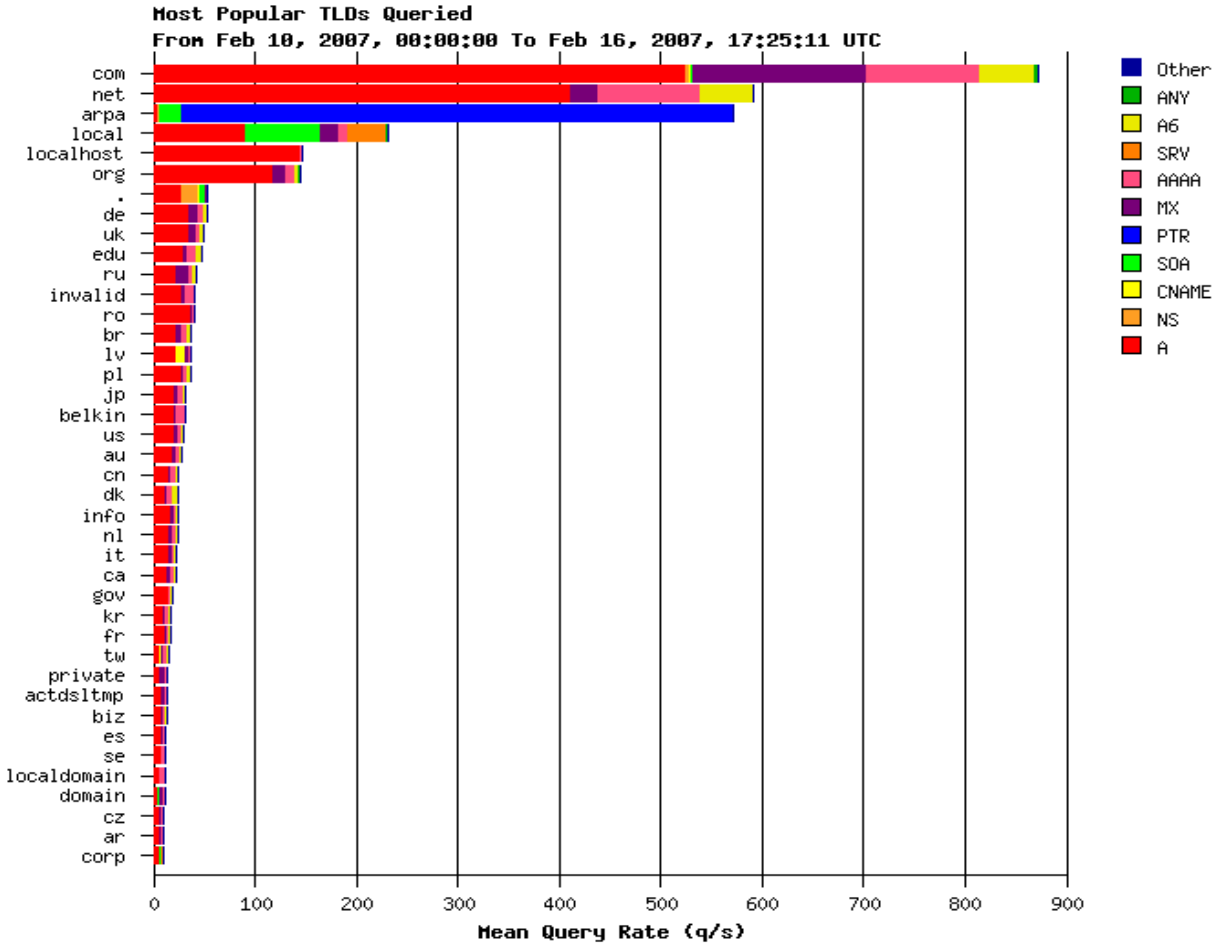

Taiwan earthquake

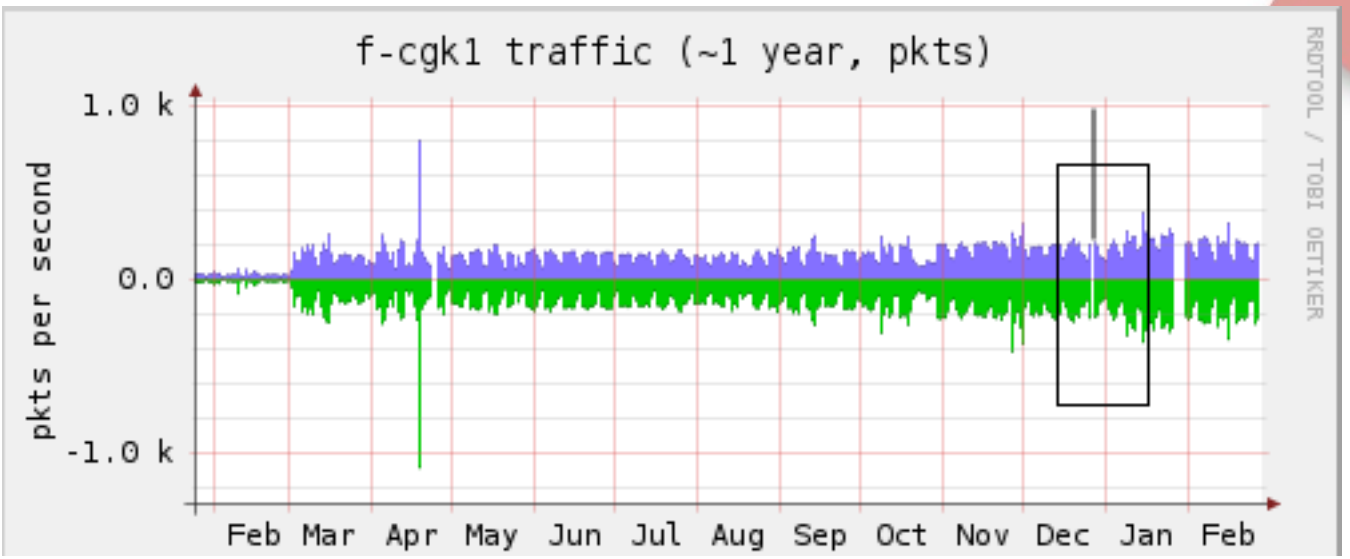

OARC Public Services

- Twice-yearly open meetings for DNS researchers and operators
  - next in Chicago 27/28th July
- <dns-operations@lists.oarci.net> mailing list
- http://public.oarci.net website
- Home for:
  - "Orphan Projects"
  - "Flood Victims"

OARC Data-Gathering Infrastructure

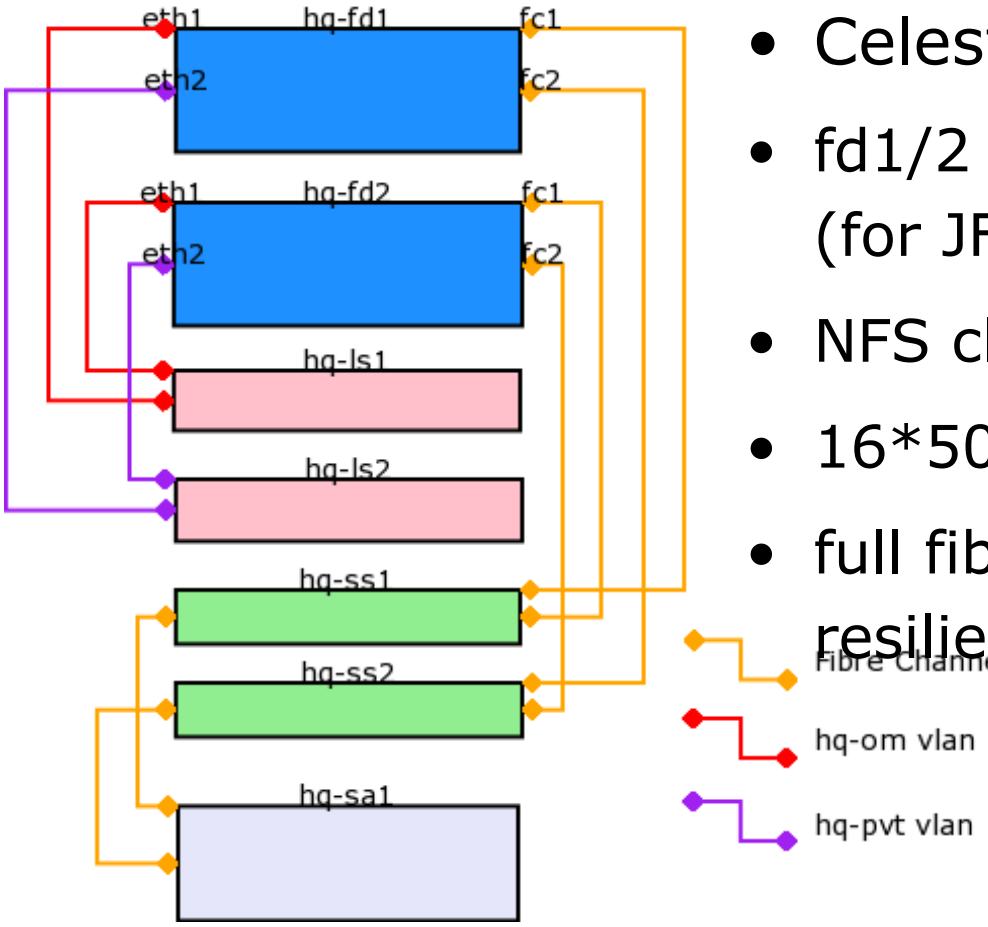

- Celestica Opteron Servers
- fd1/2 SuSE-10.1 Linux-based (for JFS support)
- NFS clients FreeBSD-based
- 16*500Gb SATA in RAID6
- full fiberchannel multipath resilience planned

A "Day in the Life of the Internet" (DITL) 8-10th Jan 2007

"Day in the Life of the Internet"

- Wide-ranging collaborative research project to improve "network science" by building up baseline of regular Internet measurement data over 48-hour periods
- See http://www.caida.org/projects/ditl/
- DNS data gathered via OARC is one part of this

DITL 8-10th Jan 2007

- OARC has supported this annually since 2004
- DNS query data gathered close to participating root and TLD servers using tcpdump into “PCAP” files
- Uploaded via ssh script to central OARC RAID system
- Available to OARC members for analysis

DITL Jan 2007 Participants

c.root-servers.net	Cogent
e.root-servers.net	NASA
f.root-servers.net	ISC
k.root-servers.net	RIPE NCC
m.root-servers.net	WIDE
as112.namex.it	NaMEX
b.orsn-servers.net	FunkFeur
m.orsn-servers.net	Brave GmbH

DITL Challenges

- Too much data
  - problem of success !
  - ran out of disk space 2 hours before end
  - “in-flight” upgrade to fix this...
- Limited space on collecting servers
- Bandwidth loss due to Taiwan quake
- Too close to seasonal holiday
- Bleeding-edge platforms

DITL Lessons Learned

- Do pending upgrades and estimate of data volumes **before** you start !
- Simple legalities = enlarged participation ☺
- Data uploading was harder than gathering
  - dry-runs helpful
- Disable auto-rotation
- Generate, preserve, share and validate data MD5 checksums
- Upgraded hardware performed well overall

DITL Results

- OARC RAID now holds over 2TB of data
  - available for research analysis
  - space for at least as much again
- Report summarising outcomes available to participants and OARC members
- More roots interested for next time
- Left us in great shape to do it again without notice 4 weeks later...

Root Server DDoS Attack
6th Feb 2007

Anycast DNS Deployment

- DNS architecture means there can only be 13 root-server IP *addresses*
- Creates potentially vulnerable bottleneck
- Anycast allows each IP address to have multiple server *instances*
- Servers geographically distributed to spread query and attack traffic
  - end-users transparently use “nearest” server
- ISC's f-root instance pioneered Anycast

F-root Anycast Instances

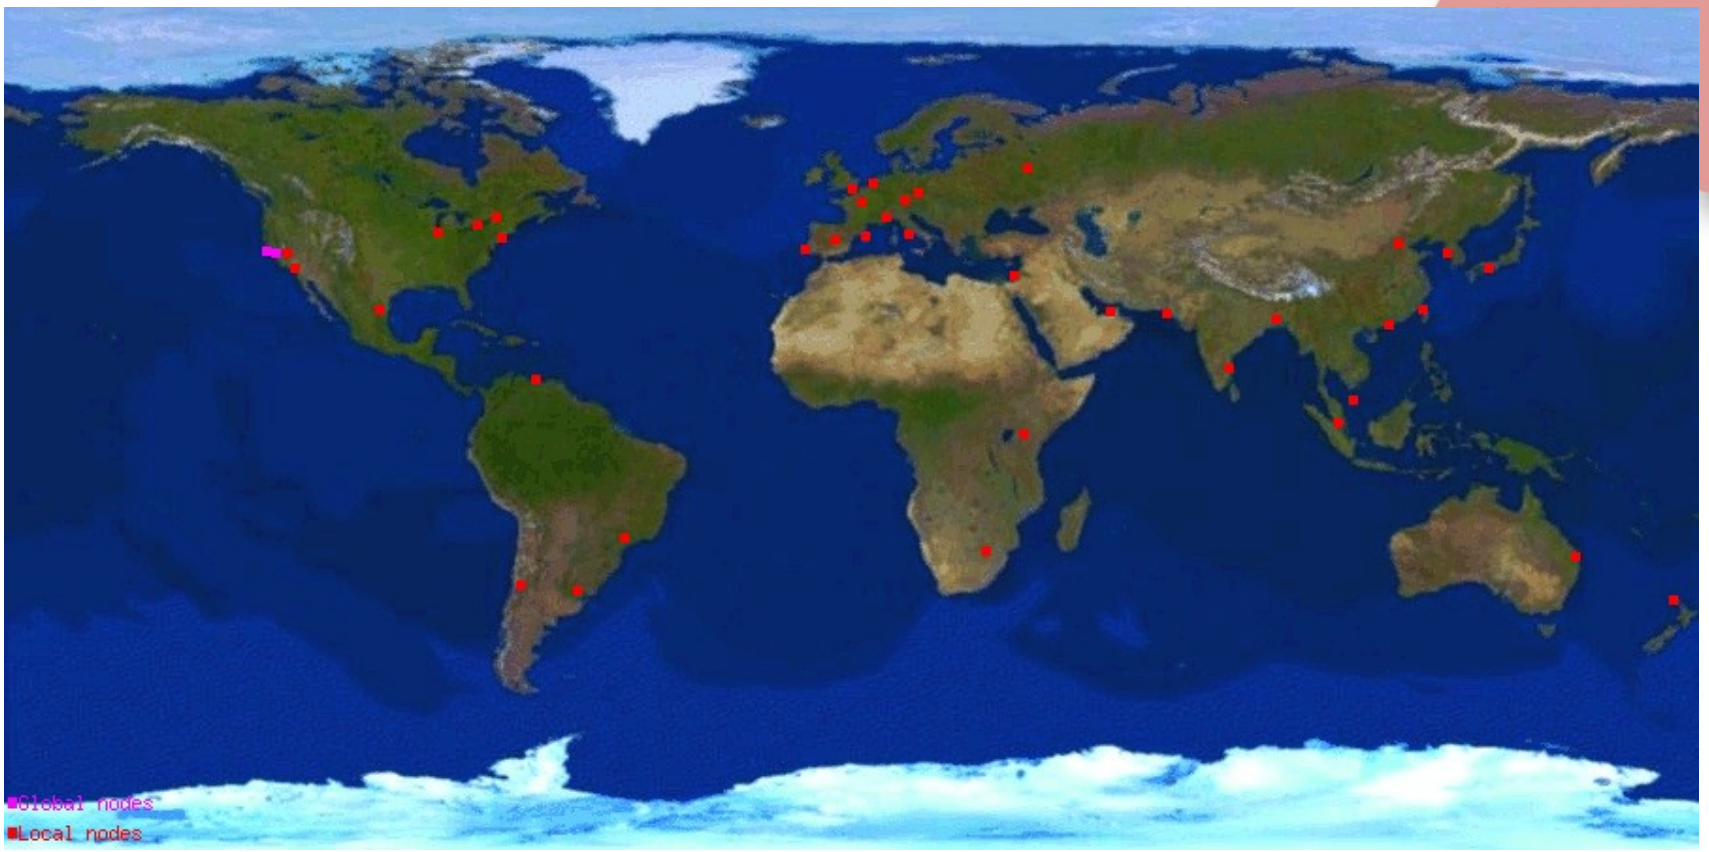

Root DDoS Attack

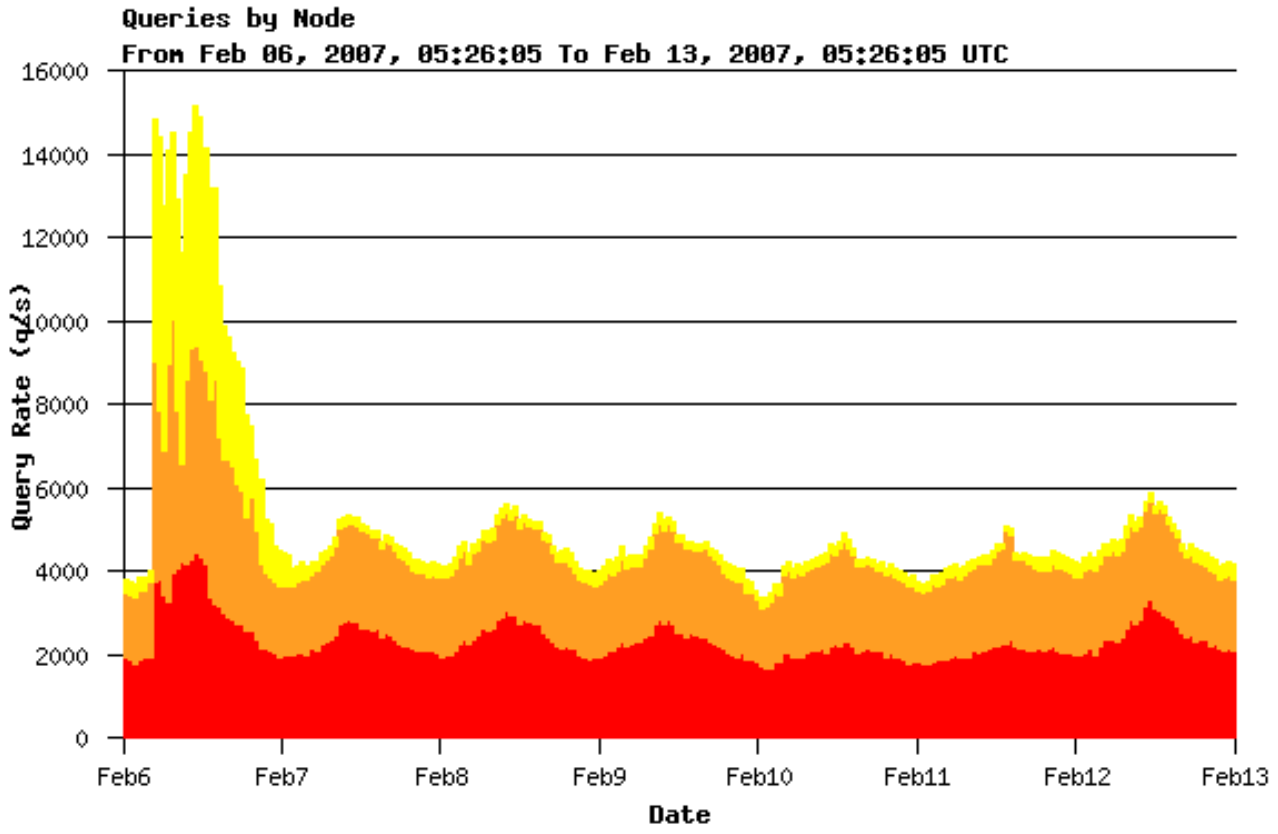

Attack overview

- Commenced at 10:00 UTC on Tue 6th Jan for 24 hours
- At least 6 Internet root and 1 TLD name servers sustained a DDoS attack
- Attack did not impact on end-user service, but was measured
- Preliminary observations made at F-root include:
  - type, quantity and distribution of attack traffic
  - how it coped
- See also ICANN report:
  - http://www.icann.org/announcements/factsheet-dns-attack-08mar07.pdf

Attack points of interest

- Happened **exactly** 4 weeks after 2007 DITL
  - may allow baseline comparison
- Happened during NANOG meeting
  - usual suspects on-hand...
- Did not use any exotic amplification techniques
- Mostly did not spoof source addresses

http://dnsmon.ripe.net

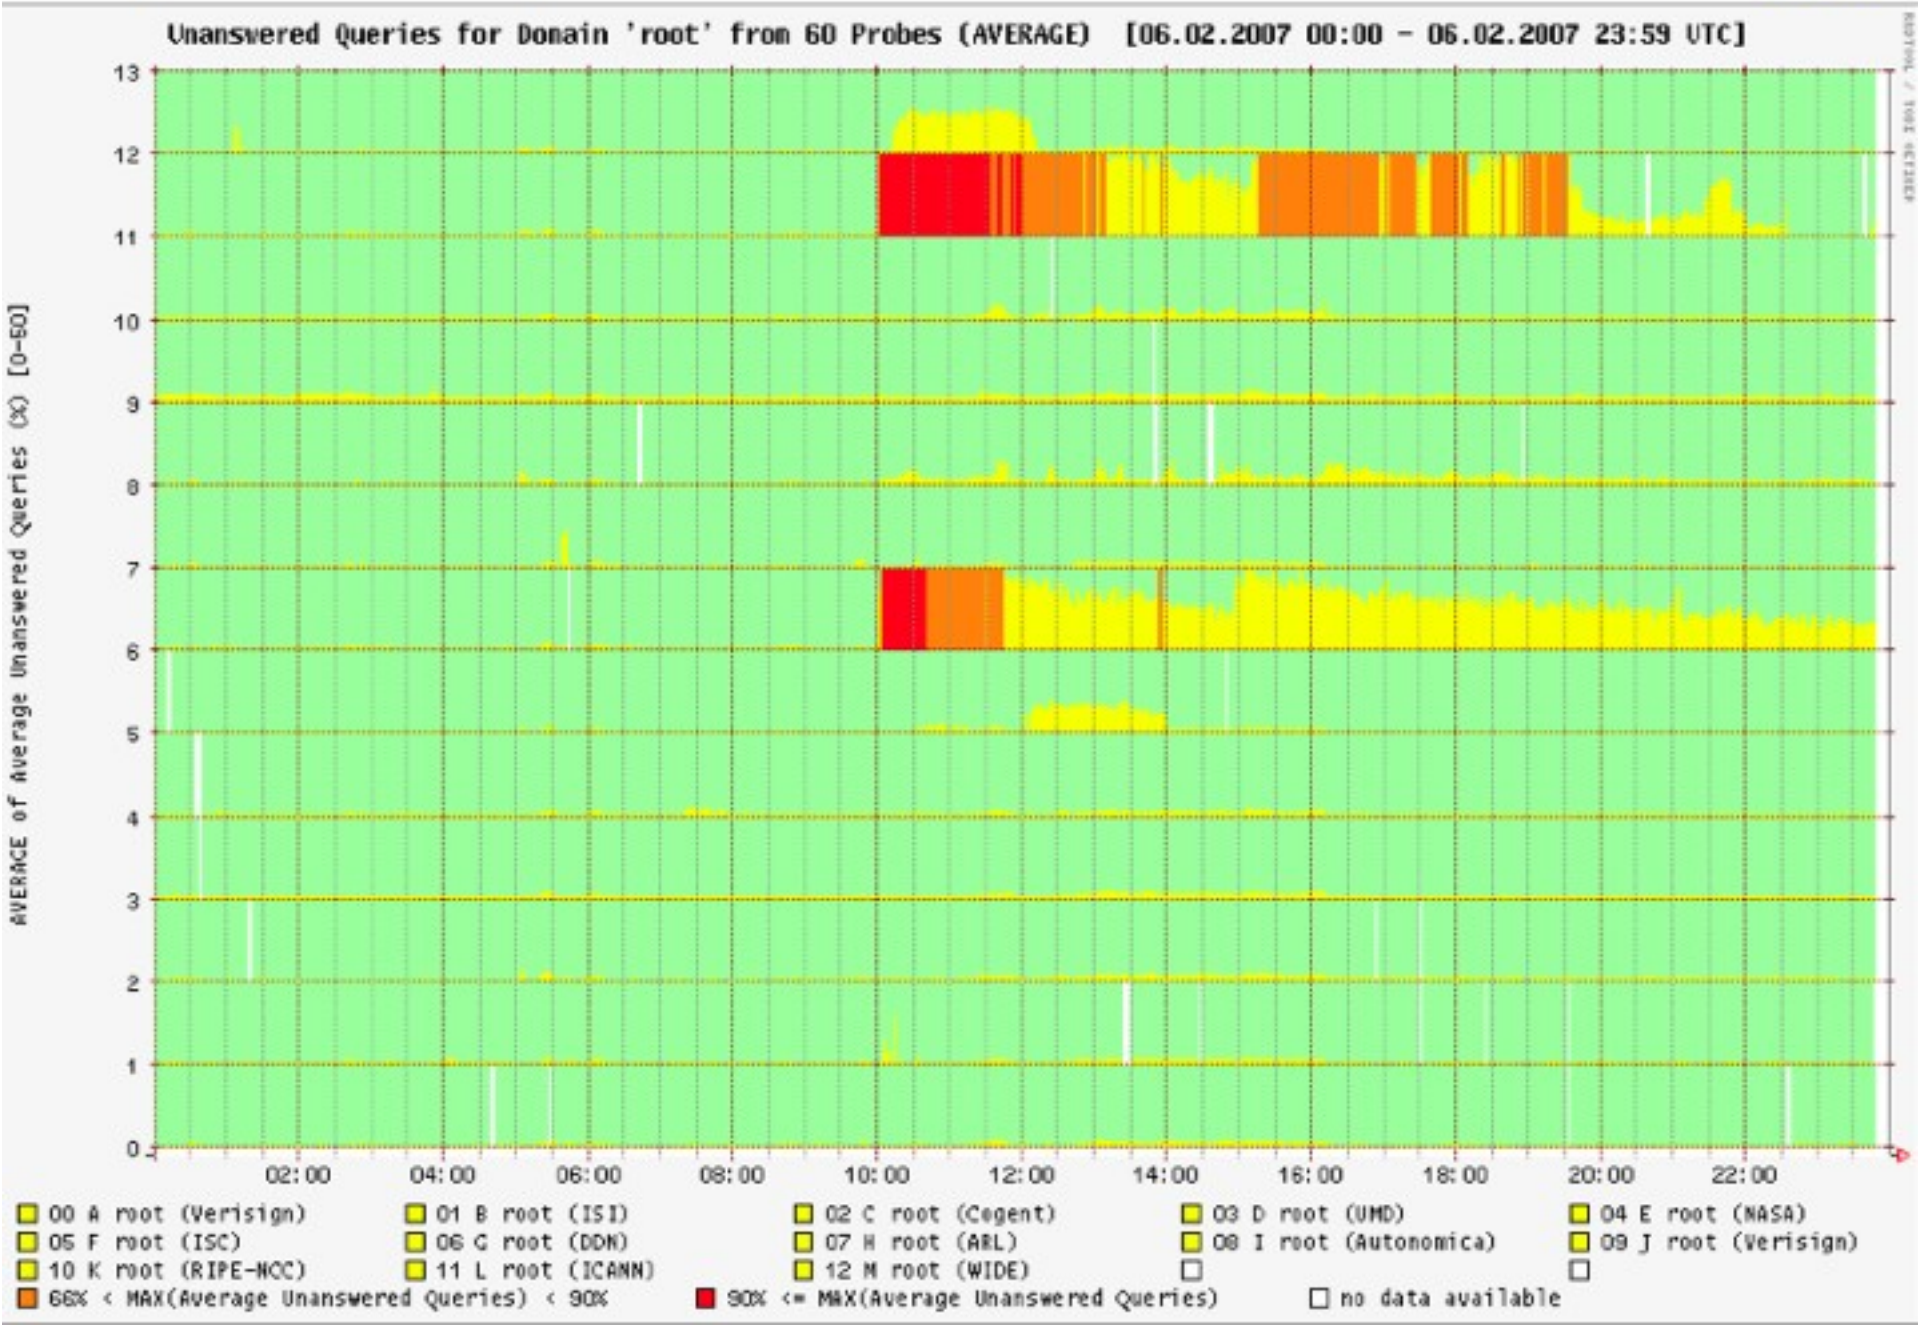

Aggregated traffic on F root

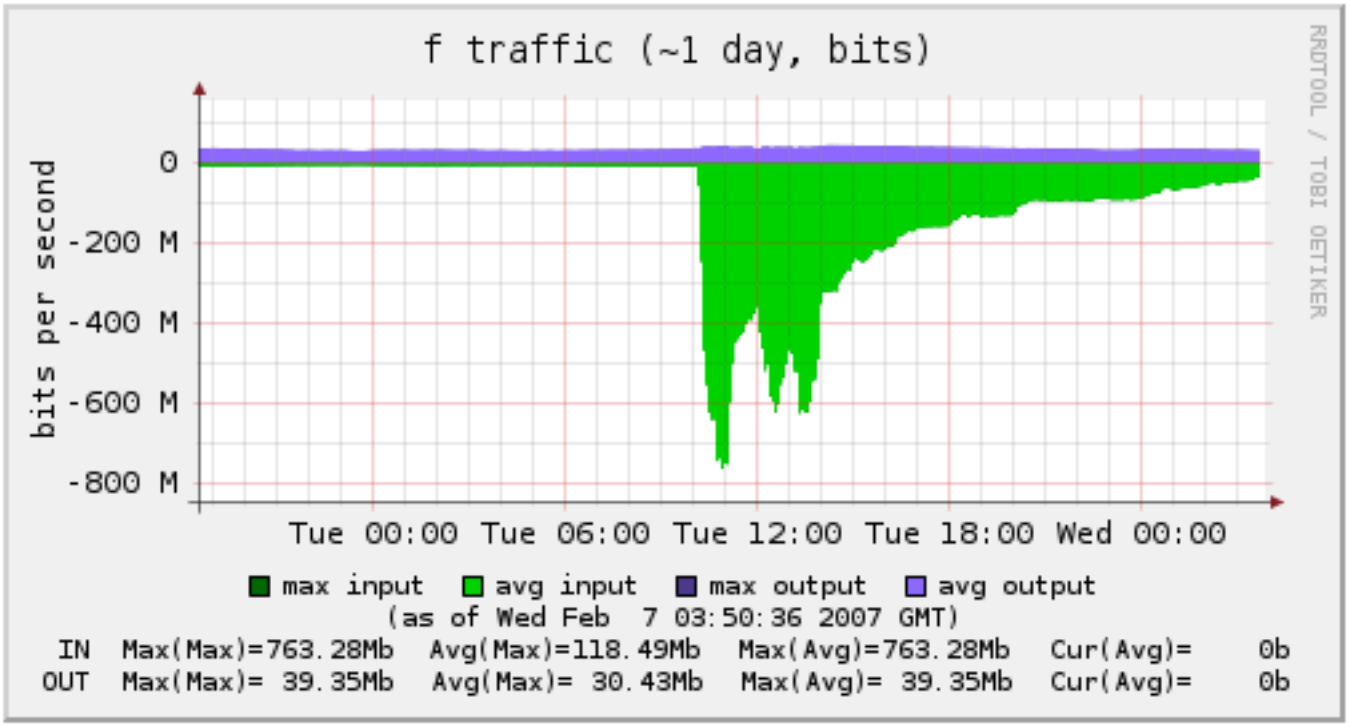

Seoul - capped at 1Gb/s

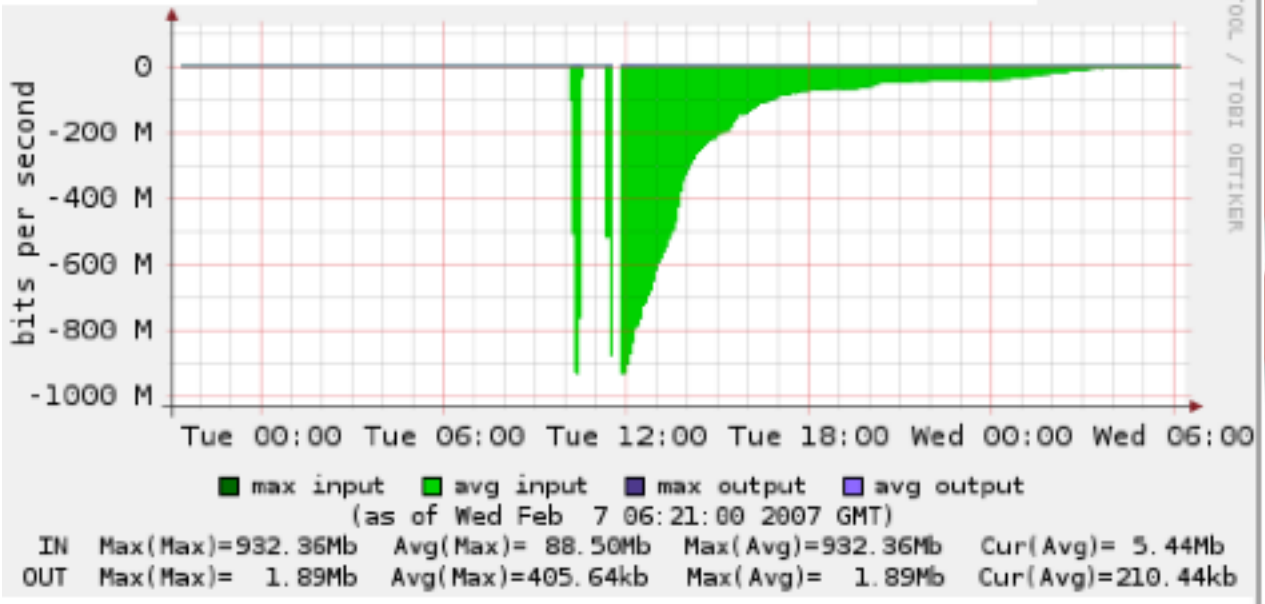

Beijing - peaked at 300Mb/s

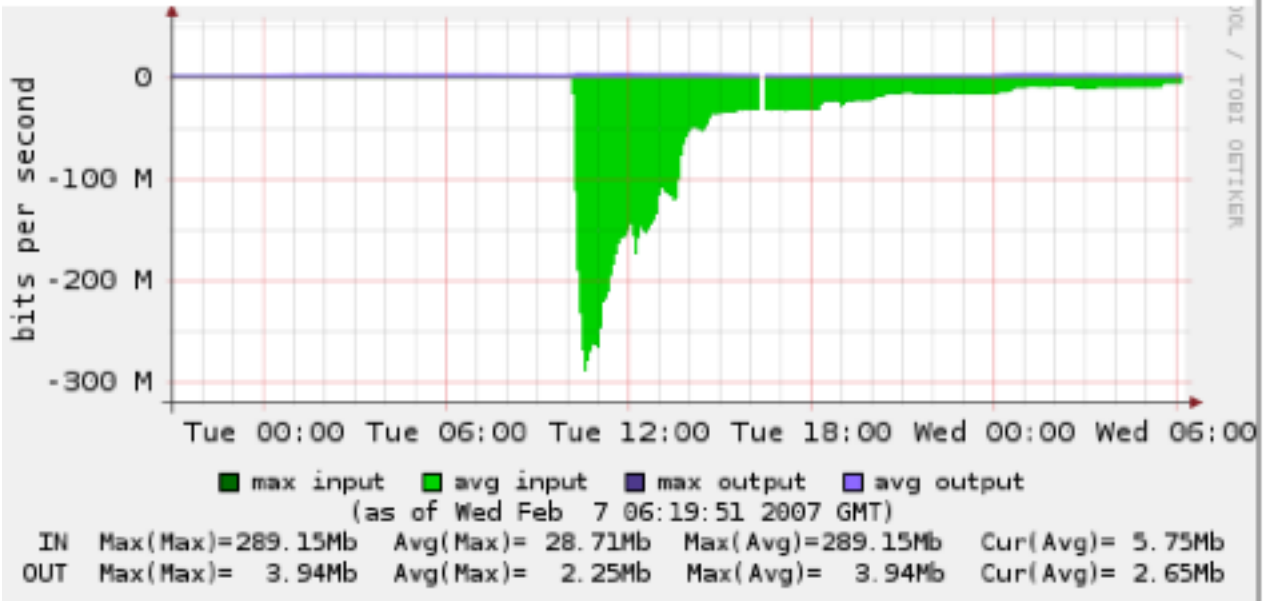

Service impact

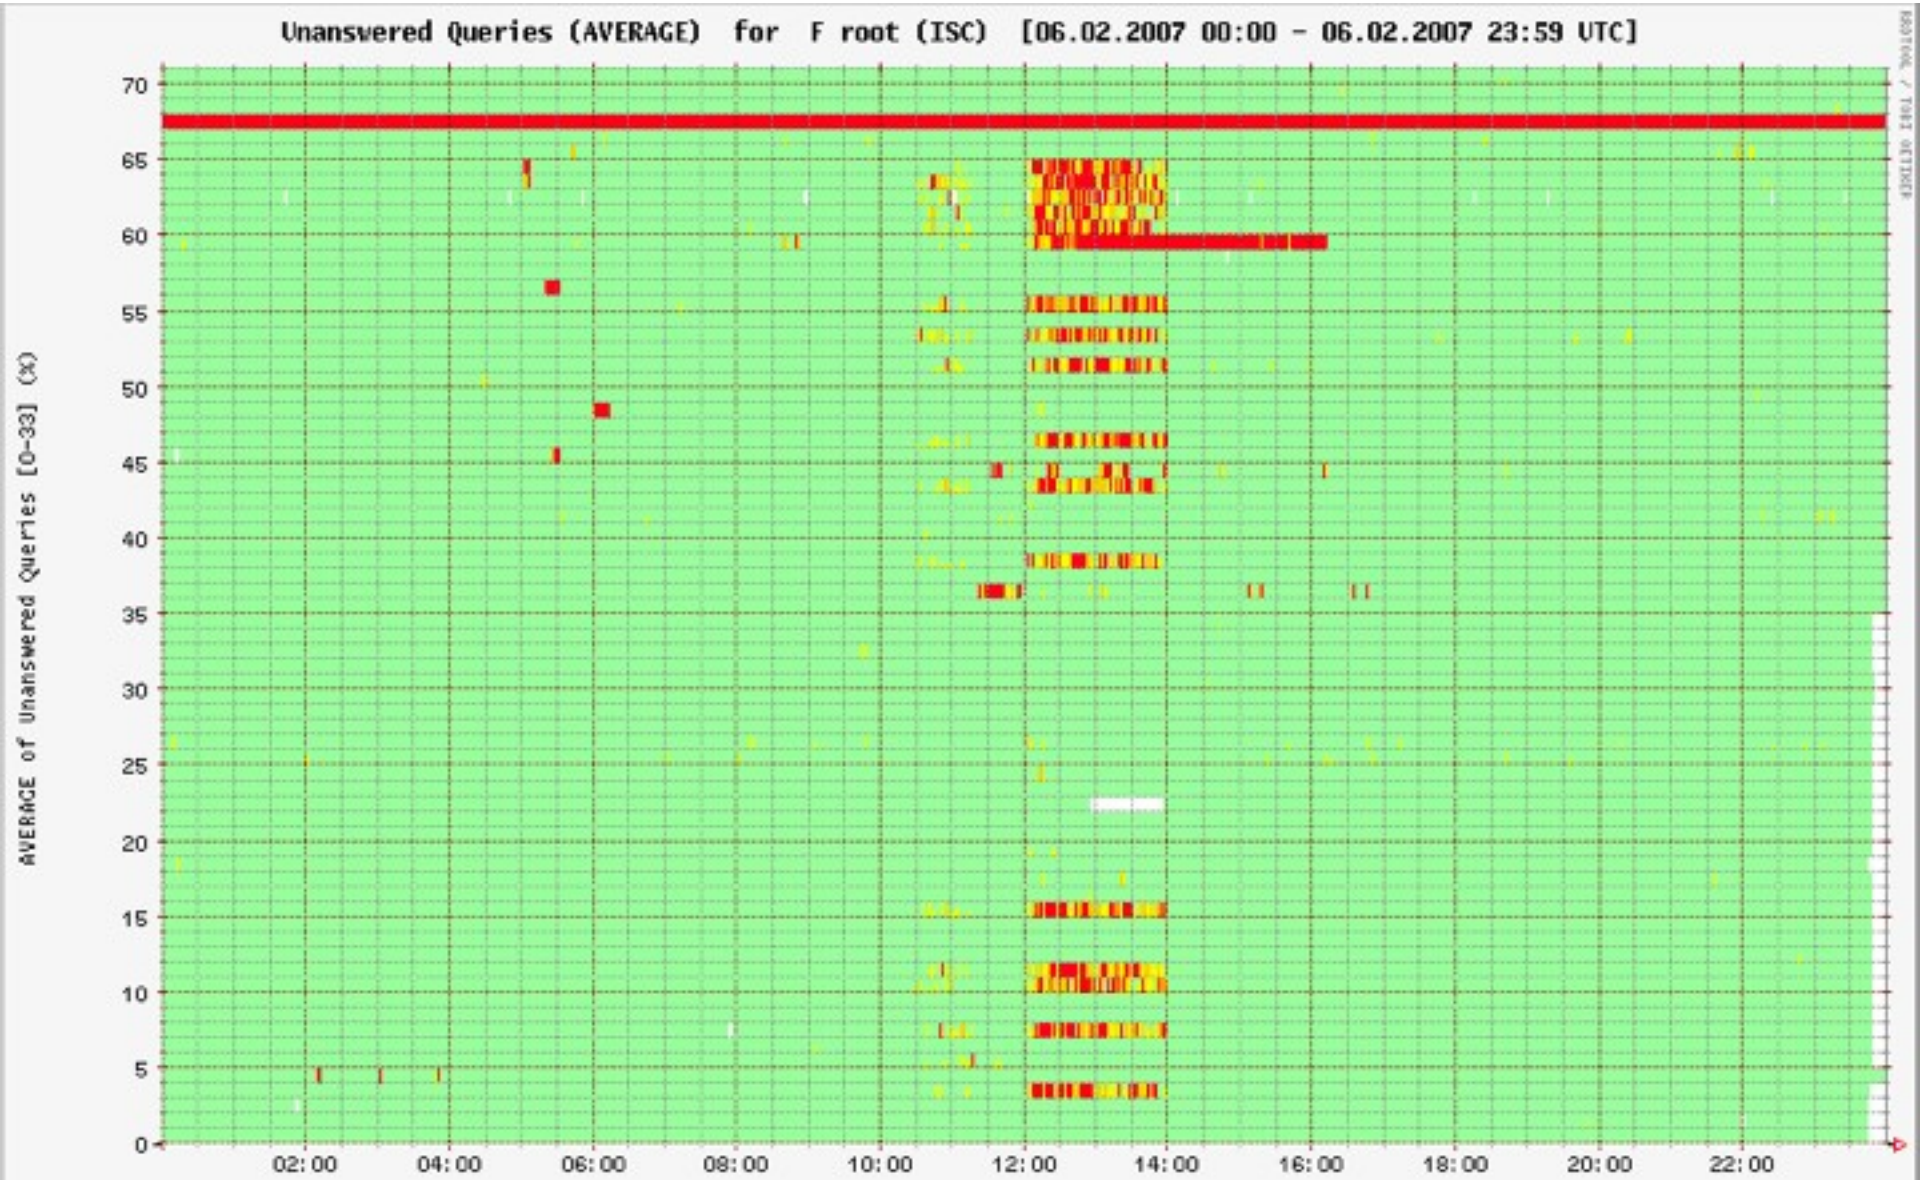

Some nodes got nothing

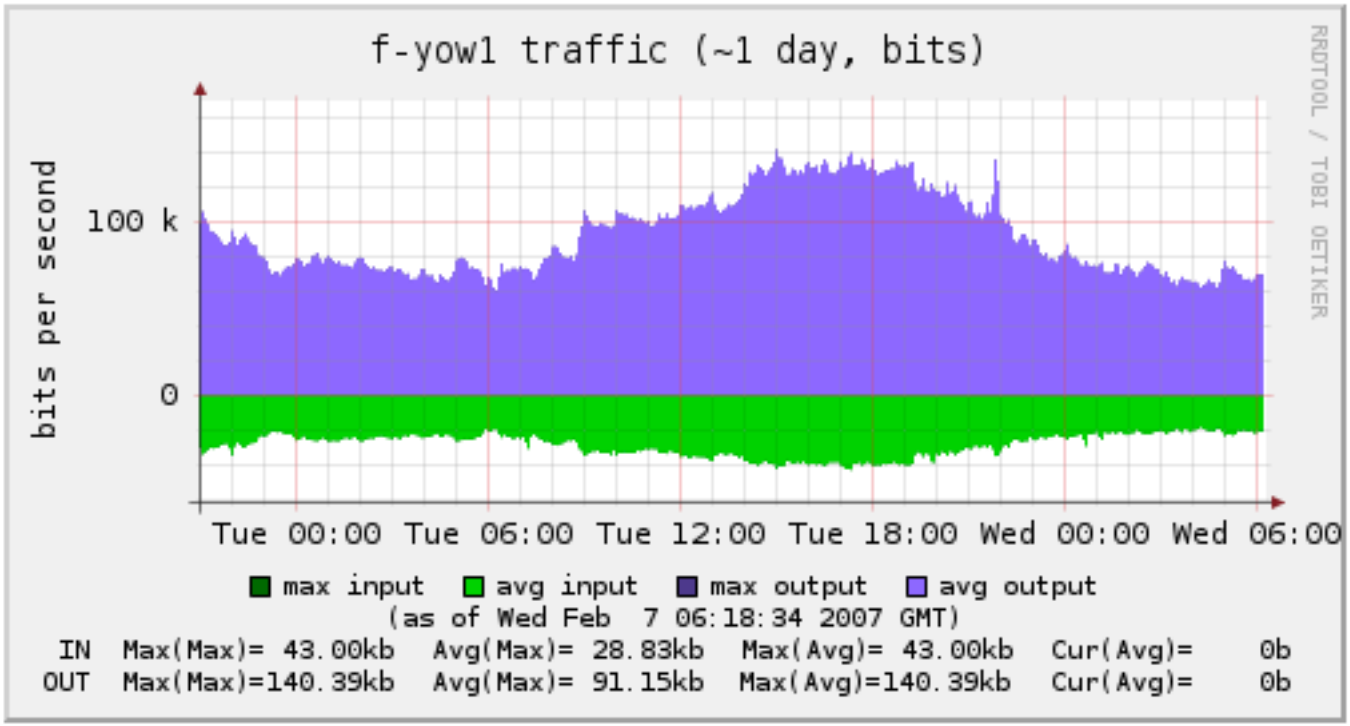

Others saw peculiar patterns

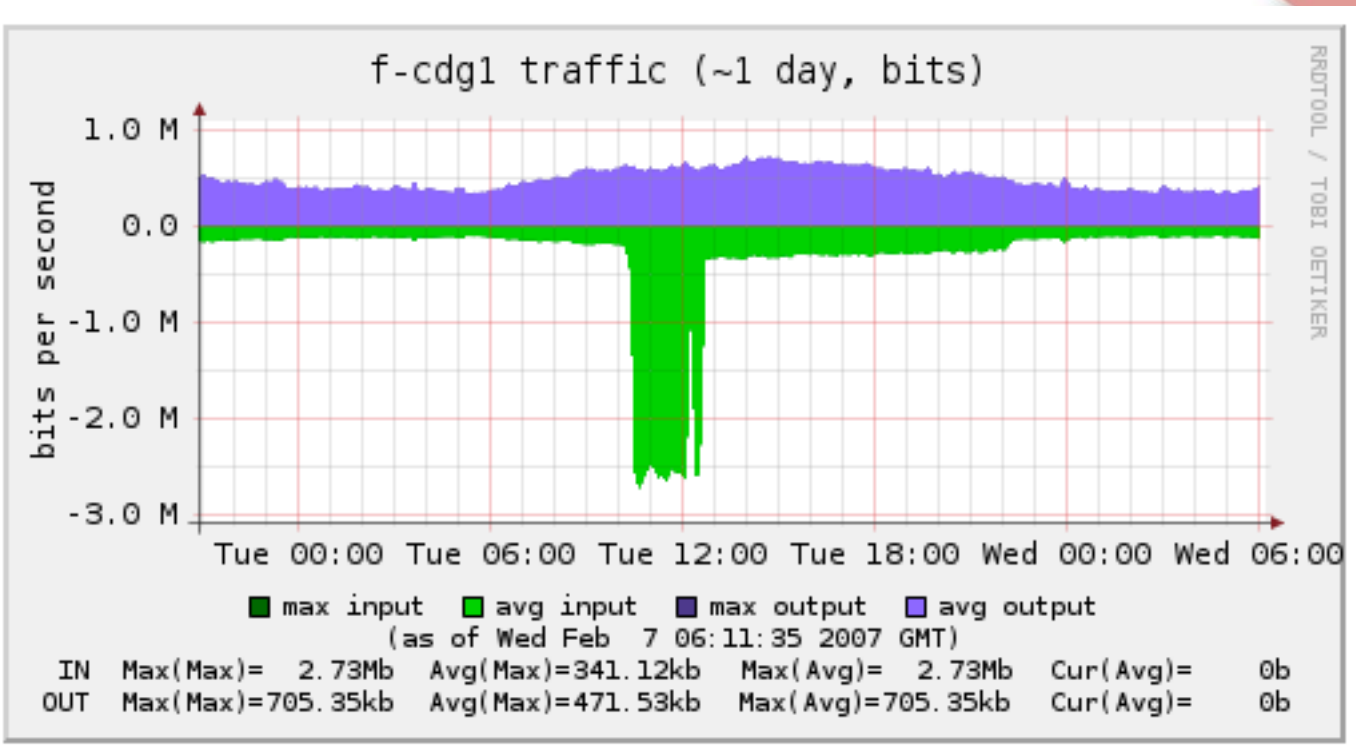

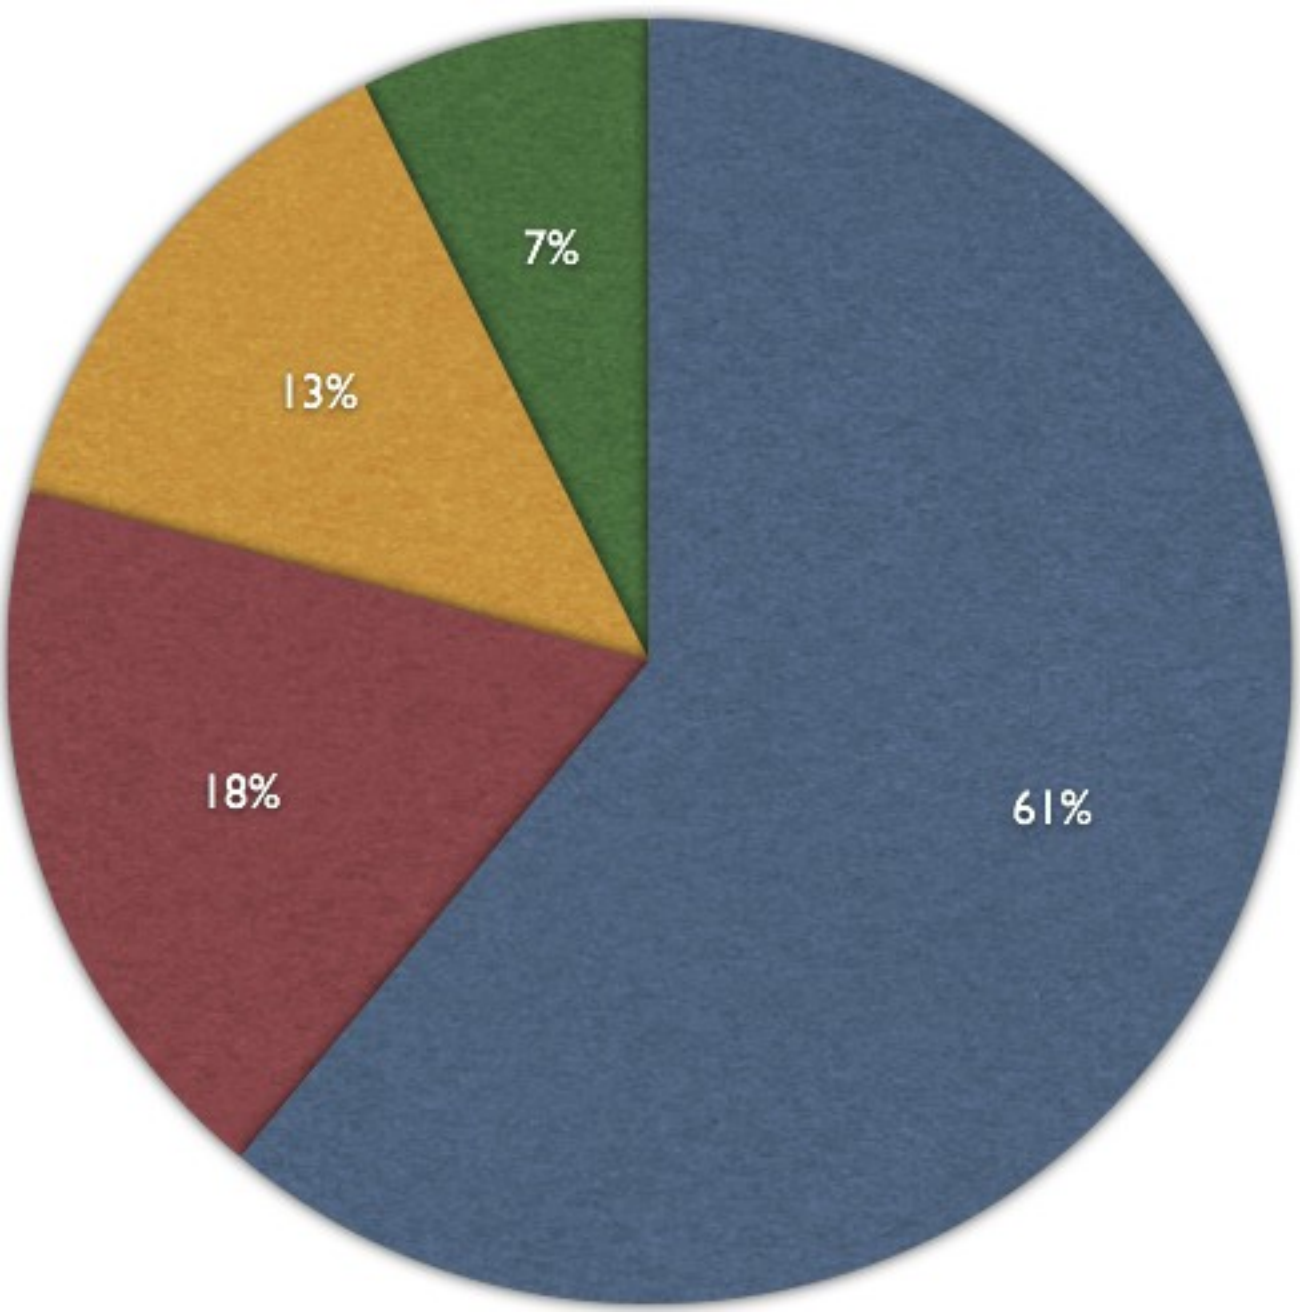

Seoul
Beijing
San Francisco
Other

Other equates to 35 F-root anycast nodes

Packet analysis

- All UDP port 53 DNS queries, containing random data
- Average size was bigger than normal traffic
  - Size random up to 1024 bytes
  - Most were more than 350 bytes
- Some were malformed DNS messages
- Contained random QTYPEs
  - updates, unknown, etc

Attack Observations

- Anycast works !
  - end-users not really impacted
  - some F-root nodes impacted, but service overall maintained
  - non-anycast nodes (G, L) hit hardest
- Filtering packets >512 bytes only partially effective
- Main sources S Korea and BellSouth, but .kr caused most of the pain
- More analysis required, will be presented at upcoming NANOG and OARC meetings

OARC Futures

- Additional resources required
- Further develop trusted communications model
- Member and open DNS Operations meeting at Chicago IETF meeting end July
- “Passive DNS” - major project to aggregate live DNS resolver data
  - seeking infrastructure funding and partners

Passive DNS

- Florian Weimer invented this concept
- Gather data close to end-user resolver servers
- Implementation in academia
  - Sensors in European ISPs & Universities
- Used today by world wide LEO community
- "Inverse directory" & botnet hunting
  - what names map to "this" address?
  - when was "this" name first used and by whom?
  - who has looked up "this" botnet C&C name?

Acknowledgments

- Dave Knight, ISC/Afilias
- Joao Damas, ISC
- John Kristoff, UltraDNS
- ICANN L-root team
- All DITL contributors

OARC Further Info

- Web: https://oarc.isc.org
- E-mail: keith_mitchell@isc.org
- Jabber: keith@jabber.oarc.isc.org
- Phone: +1 650 423 1348 (EST)
  +44 778 534 6152
- Paper:
  http://public.oarci.net/files/oarc-briefing.pdf

Questions ?