IPV6 SERVICES DEPLOYMENT

LINX IPv6 Technical Workshop - March 2009

Jaco Engelbrecht
Group Platforms Manager, clara.net

clara.net

DNS root zone goes AAAA!

- On 4th February 2008 IANA added AAAA records for the A, F, H, J, K and M authoritative name servers for the DNS root zone
- These records provide access to the root servers over IPv6 transport
- There were no significant change!

- A few weeks before that … “Hey, we need to get ns0.clara.net IPv6 enabled!”

clara.net

Before that day …

- If you wanted access to the IPv6 root zone
  - Convince Bill to give you a copy of the NDA ;-)
  - Print a copy of the NDA document
  - Sign it, fax it back to Bill

```
;       This file holds the information on root name servers needed to
;       initialize cache of Internet domain name servers
;       (e.g. reference this file in the "cache  .  <file>"
;       configuration file of BIND domain name servers).
;
;       It has both IPv4 and IPv6 capable servers and is
;       intended to test the capabilities of IPv6 native
;       transport as well as DNSSEC capabilities.

  . . . .

;       This file is made available by Bill Manning
```

clara.net

IPv6 allocation

- We use Anycast for ISP type services (Authoritative and caching DNS, NTP, etc.)
- Assigned:
  - 2001:A88:0:FFFA::/64 for IPv6 Anycast
  - a /64 for each of our Services LANs in Global Switch, Telecity and Telehouse

- Internal Anycast part of our PA – advertised to peers globally

clara.net

Enabling IPv6 on Services routers

- At the time I ran Cisco 3750s for one of our Service networks:
  - IPv6 halves TCAM available for IPv4 functions
  - No IPv6 BGP
  - No IPv6 ISIS
  - No IPv6 FHRP
- Still the case today – not ideal for v6-enabled ISP platforms

clara.net

Server Infrastructure

- We deploy /64 server LANs
- No stateless auto configuration (disabling router advertisements)
- Not using any FHRP yet (HSRP6, not great)
- Configure IPv6 address by hand
- Match last octet of IPv6 with IPv4
  - `213.253.1.24`
  - `2001:a88:0:FFF7::24`

clara.net

Authoritative DNS - PowerDNS

- One liner to enable IPv6:

```
local-ipv6=2001:a88:0:fffa::1, \
           2001:a88:0:fffa::2, \
           2001:a88:0:fffa::3
```

- Zone transfers only functional in IPv4 transport

- Since the Cisco 3750s didn't support BGP, I had to live with static routes
- Today, Quagga advertises the IPv6 Anycast prefixes via BGP

clara.net

Getting AAAA glue added

- We had to contact Network Solutions (our registrar for clara.net) to add IPv6 glue for our name servers

- Contacted Network Solutions by email
- Contacted Network Solutions by phone

  - What is IPv6?
  - What is a glue record? (!!!)
  - Escalation route did not work –

    “will take a little bit longer than usual to resolve…”

clara.net

Getting AAAA glue added ..

- Got in touch with an ex-colleague – contact in ICANN
- He exchanged a few mails, eventually got me in touch with someone ‘technical’ from Network Solutions

- When asked what the official process was:

> “In the future, please send requests to ipv6req@networksolutions.com. This is a temporary work around until this functionality is built into our account manager later this year. This information will be made public shortly via the mailing lists and also on our web page.”

clara.net

Getting AAAA glue added ...

- Added AAAA records for ns2.clara.net first
  - Waited 2 weeks to ensure no problems
- IANA added AAAA records
- Added AAAA records for ns0 & ns1.clara.net
- See http://www.sixxs.net/faq/dns/?faq=ipv6glue

clara.net

Authoritative DNS - IPv4 / IPv6 qps

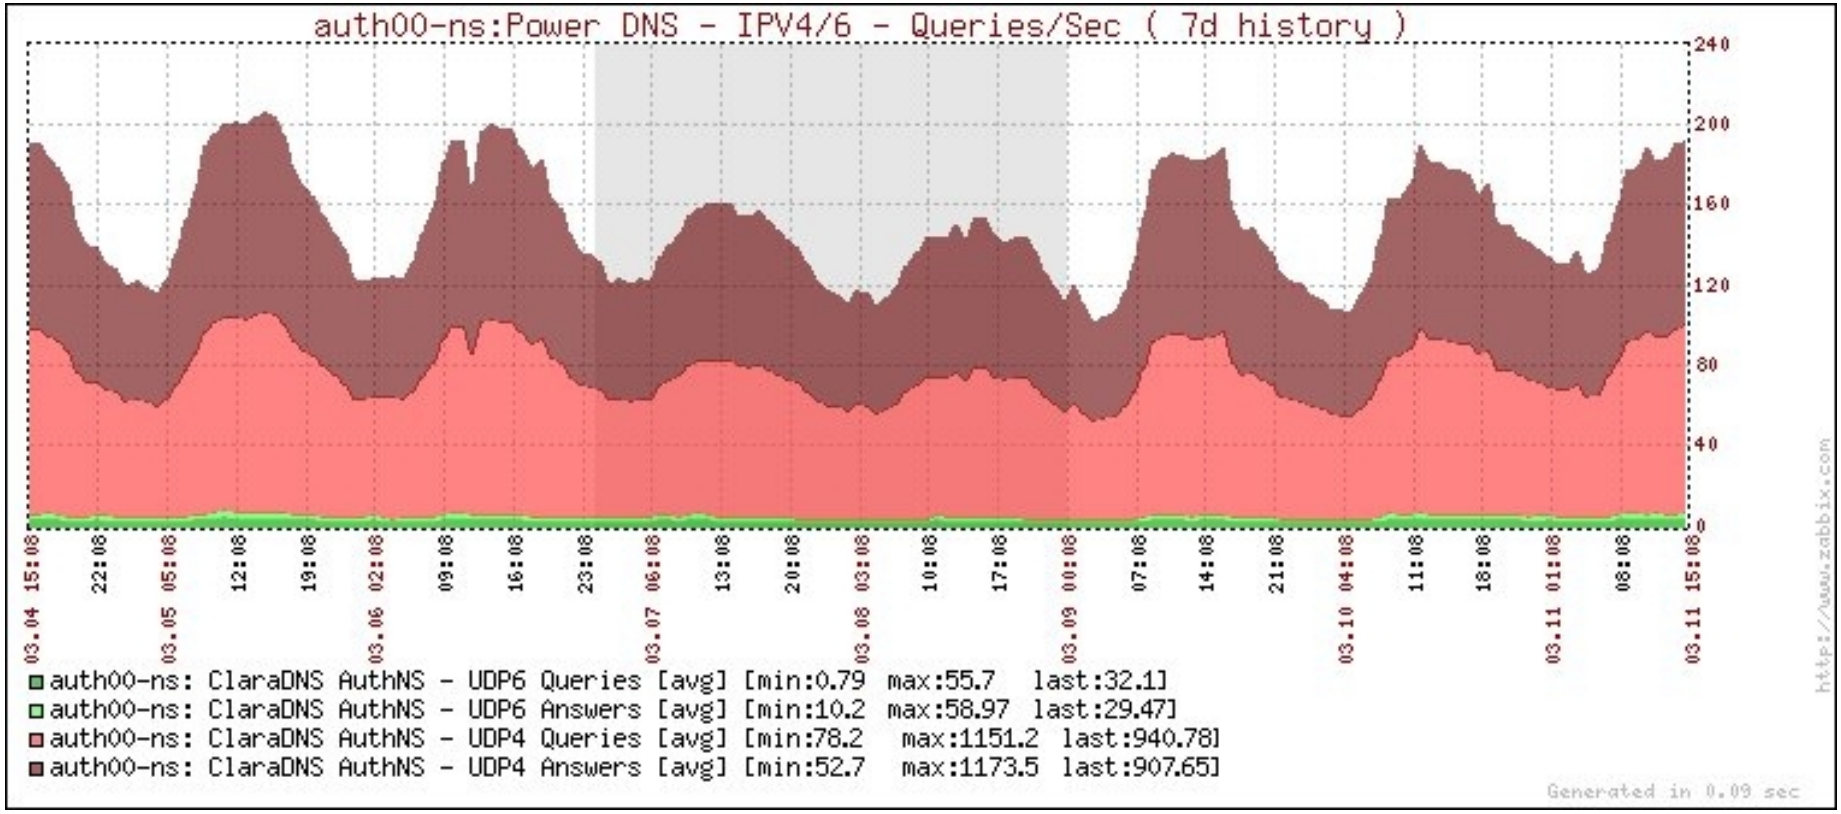

clara.net

IPV6 for public web services

- Started with noc.eu.clara.net
- Looked for reports of brokenness

- Eventually moved onto www.clara.net
  - Anycast IP, more about that and DAD later
- We have some basic traffic stats, available on noc.eu.clara.net
- Constantly looking for brokenness and ways to get better reporting
- Someone has to do it ;-)

Recursion - Unbound

- Had to update root hints file

- Resolvers –
  - Allow recursion for our IPv6 allocation
  - Performs DNS resolution via both IPv4 and IPv6

```
interface-automatic: yes
do-ip6: yes
access-control: 2001:a88::/32 allow
```

- 25% of queries are for AAAA, compared to A
- See ~4 A6 queries a day – legacy implementations

clara.net

Recursion

- RFC3901 - DNS IPv6 Transport Operational Guidelines
  - BCP for operating DNS in a v4 and v6 world

- White listing:
  - Wikimedia
  - Google – we'll be implementing this soon!

clara.net

Recursion – A, AAAA, A6 qps

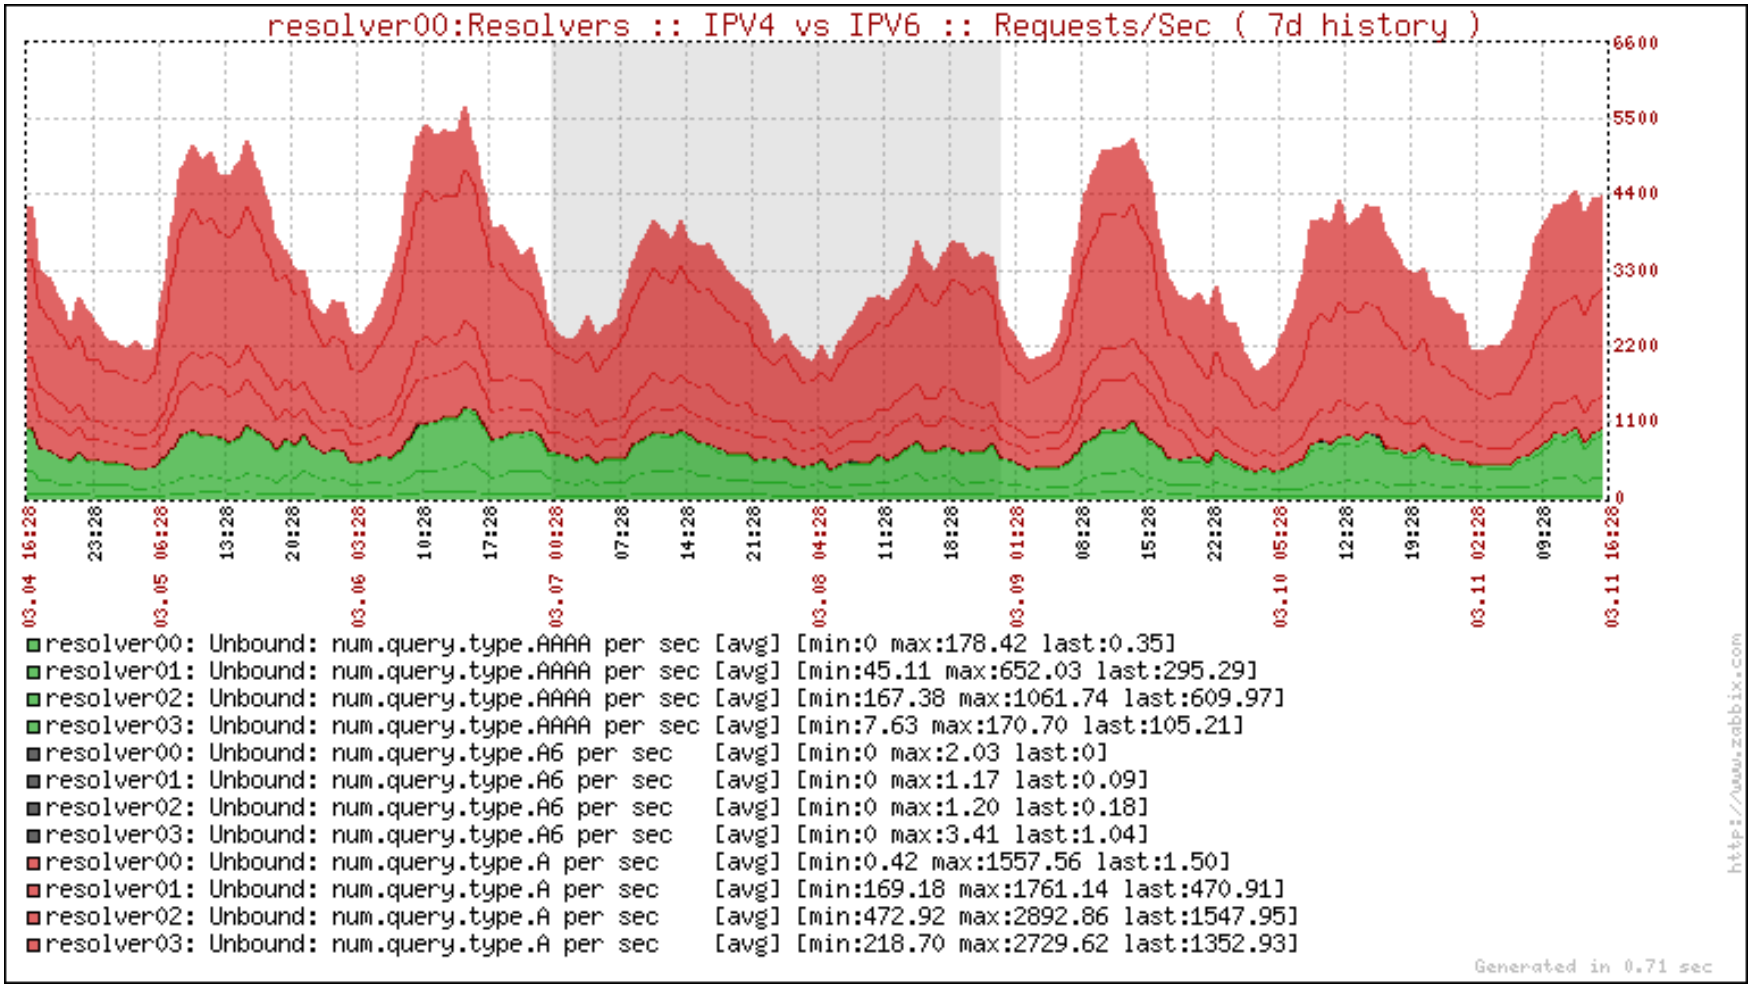

clara.net

Email

- Allow IPv6 subnets to relay mail
- Change MTA to listen on IPv6 IP address

- Blacklists
  - No operational IPv6 blacklists yet
  - A few test ones floating around, none that we'd use in production

- Spamhaus mirror IPv6 enabled

clara.net

bje@mana:~$ dig +short j.ns.spamhaus.org. AAAA

Email - Exim

- Changed Exim to only do RBL lookups if it's IPv4 sender:

```
deny
        condition = ${if isip4{$sender_host_address}{yes}{no}}


         dnslists  = some.rbl
```

clara.net

Syslog, NTP, Courier-IMAP, Apache, FreeRADIUS

- Most of these ‘just work’ now, with full IPv6 support
- Not the case a few years ago – had to apply patches
- No official IPv6 support in MySQL yet for the network stack, nor data types

clara.net

Network infrastructure

- Management VLAN for services network infrastructure is IPv6 enabled
- All services switch, router, load balancer management can be done via either IPv4 or IPv6

- Syslog over IPv6

clara.net

Load balancers

- We use Citrix Netscaler AS 7000s
- Recently upgraded to version 9
- IPv6 support now available – for free!

- Two ways:
  - IPv6 Offload Mode – runs a V6-V4 lookup table – translates an IPv6 request into a V4 request to backend
  - IPv6 end-to-end

clara.net

Monitoring

- Nagios
  - Custom plugins written in Perl to perform IPv6 tests
- Zabbix 1.6
  - All modules support both IPv4 and IPv6
- Use Zabbix as primary monitoring tool
  - Planned upgrade to 1.6 later this quarter
  - Will migrate Nagios IPv6 monitoring over then

clara.net

Securing it all - routers

- Router ACLs generated by internal tool
- This was patched to support IPv6
- Strange stuff going on in IOS:

```
permit tcp any host 2001:DB8:0:B33F::13 eq 995
permit tcp any host 2001:DB8:0:B33F::13 eq 993
sequence 312 permit tcp 2001:DB8:3:B::/64 host 2001:DB8:0:B33F::1A eq
  443
sequence 313 permit udp any host 2001:DB8:0:B33F::54 eq syslog
sequence 320 deny tcp any any log
deny udp any any log
deny ipv6 any any log
```

clara.net

Securing it all - systems

- Match your existing IPv4 ACLs at the very least (or better, review and fix ‘em!)
- Nessus supports IPv6 targets
- Many software packages listens on tcp6 by default these days
  - You’ve got an IPv6 ACL to block inbound tcp/22, right? ;-)
- /etc/sysctl.conf – disable the defaults!
  - `net.ipv6.conf.all.autoconf=0`
  - `net.ipv6.conf.all.accept_ra=0`

clara.net

Anycast / NIC Teaming and DAD

- Duplicate Address Detection (DAD)
- Causes duplicate checks on team and physical NICs even though the physical is not used for addressing

```
# sysctl -w
net/ipv6/conf/bond0/dad_transmits=0
```

```
# sysctl -w
net/ipv6/conf/eth0/dad_transmits=0
```

clara.net

PXE boot

- No ratified IPv6 PXE standard yet that vendors can comply with
- Going to have to kickstart our servers with IPv4 for the time being

clara.net

Disabling IPv6 (if you need to ;)

- Debian
  - Blacklist the module - /etc/modprobe.d/blacklist
    - *blacklist ipv6*
- Firefox
  - In about:config, set network.dns.disableIPv6 to true
- Exim
  - disable_ipv6 in Global configuration

clara.net

Troubleshooting

- netcat6 (nc6)
- netstat | grep tcp6
- pchar
- tracepath6 for finding those MTU change points

- Wireshark / tshark
- ipv6calc / sipcalc
- Perl – ‘use IO::Socket::INET6’

clara.net

Mailing lists

- ipv6-ops@lists.cluenet.de
- users@ipv6.org
- debian-ipv6@lists.debian.org

clara.net

Conclusion

- Many of the challenges are superficial because the underlying architecture and principles of have not changed
- Enabling IPv6 gives us the opportunity to do things differently, better and improve on many IPv4 practices, such as:
  - IP Addressing plans
  - ACLs
  - Monitoring

clara.net