

Pre Talk

working with Akamai
to improve HTTP(S) traffic flows

Thomas Mangin – LINX 97 – 20th and 21st February
Thank you to Ronan Mullally from Akamai Technologies for his help

The players

Ronan Mullally, playing for team 
A large CDN, “generator” of “desirable” content

And “Me”, playing for team  exa networks
An Akamai peer

The game

Better, Happier Peering



Exa Networks

Education & Business ISP

Eyeball Network

10 Gb National Core

- London, Manchester, Leeds ring
- Bradford, Sheffield Q1, more locations soon

100 Gb within POPs (using Arista 7280R series)

DSL/FTTC traffic

slightly over 2/3 comes from Manchester

slightly under 1/3 comes from Telehouse Docklands

Leased Lines

“even split” London, Manchester & Leeds

Dark Fiber

Bradford & Leeds (ATM)

Exa Networks Limited

Organization	Exa Networks Limited
Also Known As	
Company Website	https://www.exa.net.uk/
Primary ASN	30740
IRR Record	AS-EXA
Route Server URL	
Looking Glass URL	
Network Type	Cable/DSL/ISP
IPv4 Prefixes	200
IPv6 Prefixes	200
Traffic Levels	5-10Gbps
Traffic Ratios	Mostly Inbound
Geographic Scope	Regional
Protocols Supported	<input checked="" type="radio"/> Unicast IPv4 <input type="radio"/> Multicast <input checked="" type="radio"/> IPv6
Last Updated	2016-06-21T11:18:33Z
Notes	We prefer to see networks we exchange little traffic with through route-servers. In general we have an open peering policy except AMS-IX, France-IX and NL-IX where we peer selectively.

Peering Policy Information

Peering Policy	http://as30740.net/
General Policy	Open
Multiple Locations	Preferred
Ratio Requirement	No
Contract Requirement	Not Required

Contact Information

Role ▾	Name	Phone E-Mail
NOC	NOC	00 44 345 145 1234 noc@exa-networks.co.uk
Policy	Thomas Mangin	00 44 345 145 1234 peering@exa-networks.co.uk
Sales	Sales	00 44 345 145 1234 sales@exa-networks.co.uk
Technical	Richard Halfpenny	00 44 345 145 1234 noc@exa-networks.co.uk
Technical	Daniel Piekacz	00 44 345 145 1234 noc@exa-networks.co.uk

Exa Networks

Most traffic inbound

- **1/3 Akamai**
- **1/3 Google**
- Apple, Microsoft, Netflix, Limelight, Amazon, Facebook
("the usual suspects", traffic volume varying depending on the week)
- The *very* long tail

Our own 'in-house' **content filtering** solution

HTTP and HTTPS traffic - proxied & trans-proxied

Bringing back some traffic from London to Manchester

We provide Akamai with some **transit** for their IXLeeds cluster
(As it will be *very* visible on our AS-Stats graphs)

Not selling AS30740 transit otherwise (but friends and family)

Very **open peering policy** (1,000+ IPv4 & IPv6 eBGP sessions)

We may need to reduce the number of sessions at some point

Exchange ▼

ASN

Speed

RS Peer

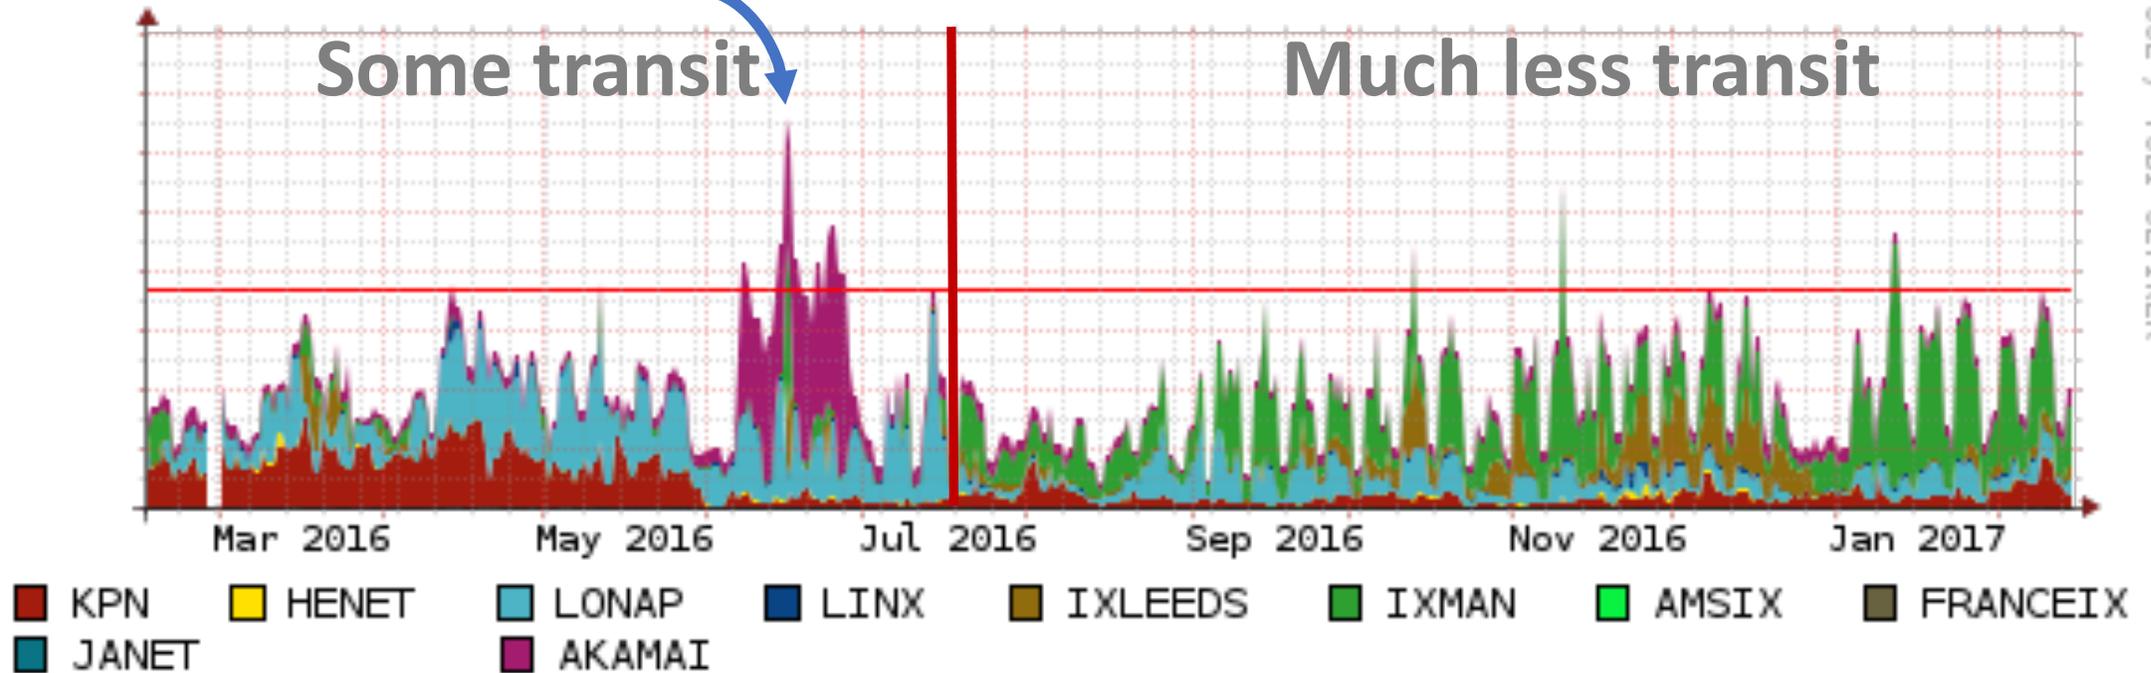
Peering sessions with Akamai

IX Leeds	10G
30740	<input type="radio"/>
IXManchester	10G
30740	<input checked="" type="checkbox"/>
LINX LON1 Main	10G
30740	<input checked="" type="checkbox"/>
LONAP LON0	10G
30740	<input checked="" type="checkbox"/>

Looking over a long period

IXLeeds' Akamai cluster sending to Asia, using the transit link we provide them

Change



RRDTOOL / TOBI OETIKER

Mostly from **LoNAP**

Mostly from **IXManchester**

Something happened during the summer 2016 !

~~BGP~~ DNS Based Routing

“Adding” a recursive DNS server per POP, we AnyCasted our DNS service IPs in:

- London
- Manchester
- Leeds

Previously the setup was Active / Passive using LocalPref.

A DNS server failure will cause another POP DNS to be used.

AnyCasting DNS

Eating our own dog food, ExaBGP is used to announce all our /32 service IPs, including for DNS.

Detect DNS failure and stop announcing the service IP should DNS fail to resolve.

Let me Google that for you:

`“exabgp healthcheck anycast DNS filetype:pdf”`

PowerDNS

And we love PowerDNS / dnsmdist. You can do some funky things with it, like use LUA to manipulate DNS answers.

Explanation: from the mouth of the network

(UK peering forum presentation)

How does it work?

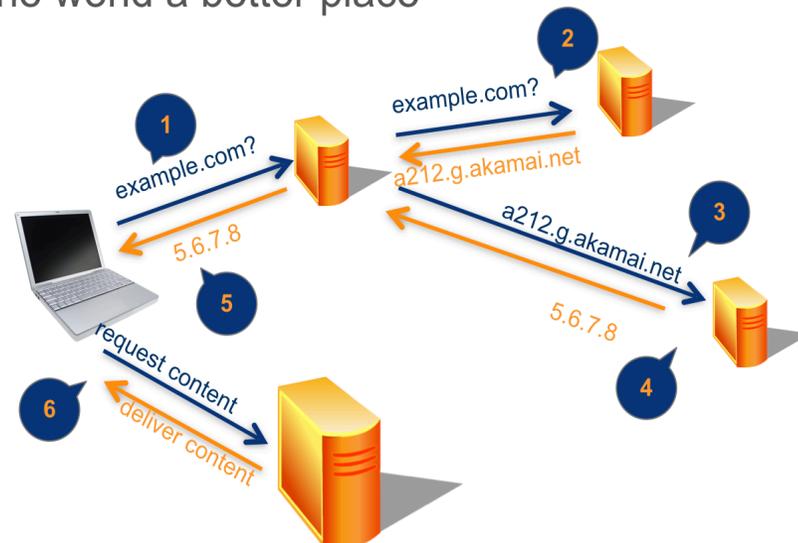
The 'secret sauce' of a CDN is its ability to direct end-users to content

We call this Mapping

Proper Mapping makes the world a better place

- better performance
- lower cost
- happy networks
- Happy Users

Akamai maps using DNS lookups



©2015 AKAMAI | FASTER FORWARD™

Akamai needs to know where your DNS are
(or they will need to guess from latency)

Your user will be served based on its DNS resolver
(DNS in London or 8.8.8.8, traffic from London)

DNS traffic level matters

If a DNS is not generating enough traffic, it will be consolidated with others.

Akamai monitor and maps your DNS servers
(path taken, latency, packet loss, ...)

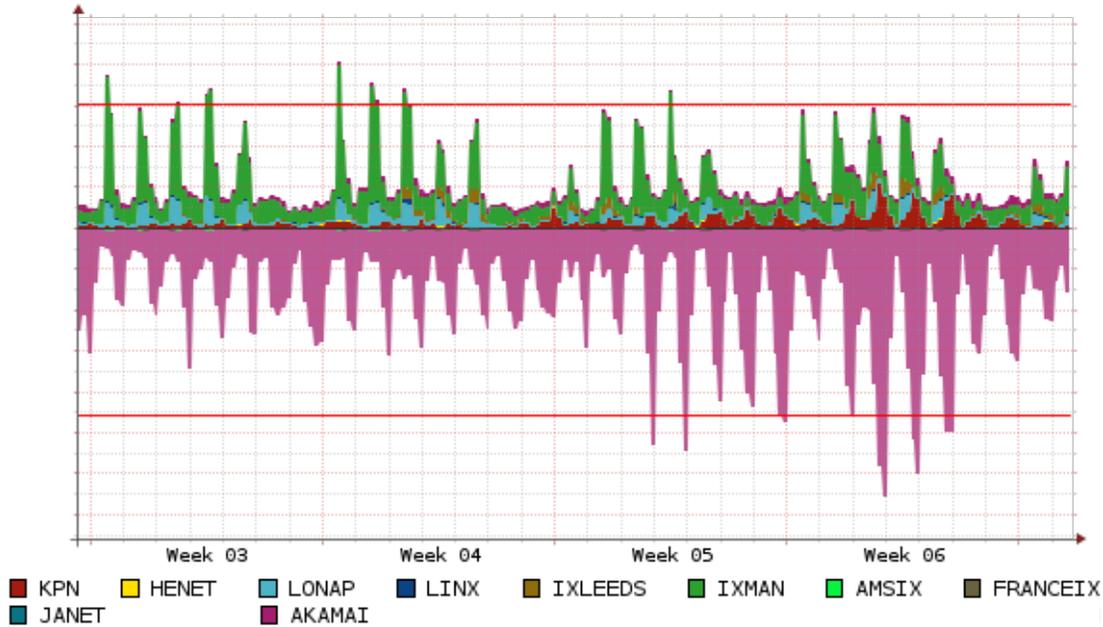
Packet loss/congestion on a link affect the decision

If you are using RFC 7871, let Akamai know, they may take advantage of it.

Client Subnet in DNS Queries

<https://tools.ietf.org/html/rfc7871>

AS20940 - Akamai International B.V. - IPV4



RRU700L / TOBI OETIKER

Traffic coming mostly from

- 1 - IXManchester (green)
- 2 - LoNAP (light blue)
- 3 - Transit in Manchester (red)
- 4 - IXLeeds (brown'ish)
- 5 - The rest

This and last weeks are “quieter school holiday” weeks but the ratio is still representative.

Trying to understand that data, let's put some price on the 10Gb links:

LON1 - £834 pcm (£424 LON2)

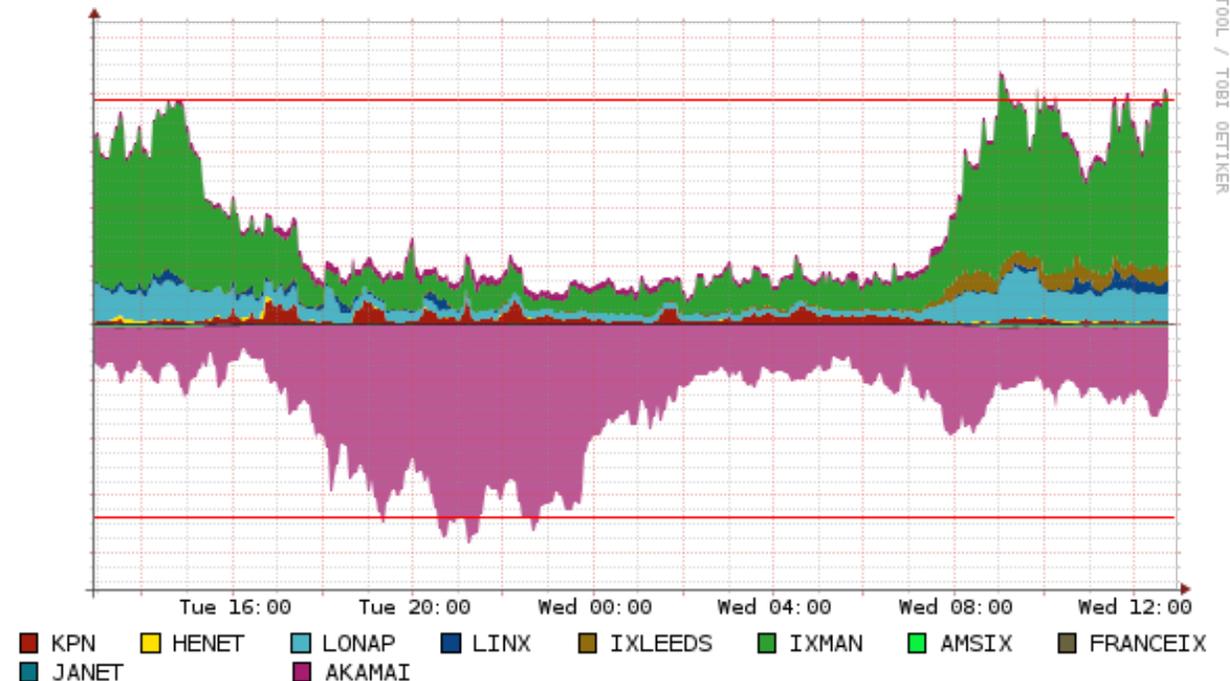
LoNAP - £375 pcm (so LoNAP cheaper in London)

IXMan - £350 pcm

Regional peering traffic level is not affected by the price of other IXes.

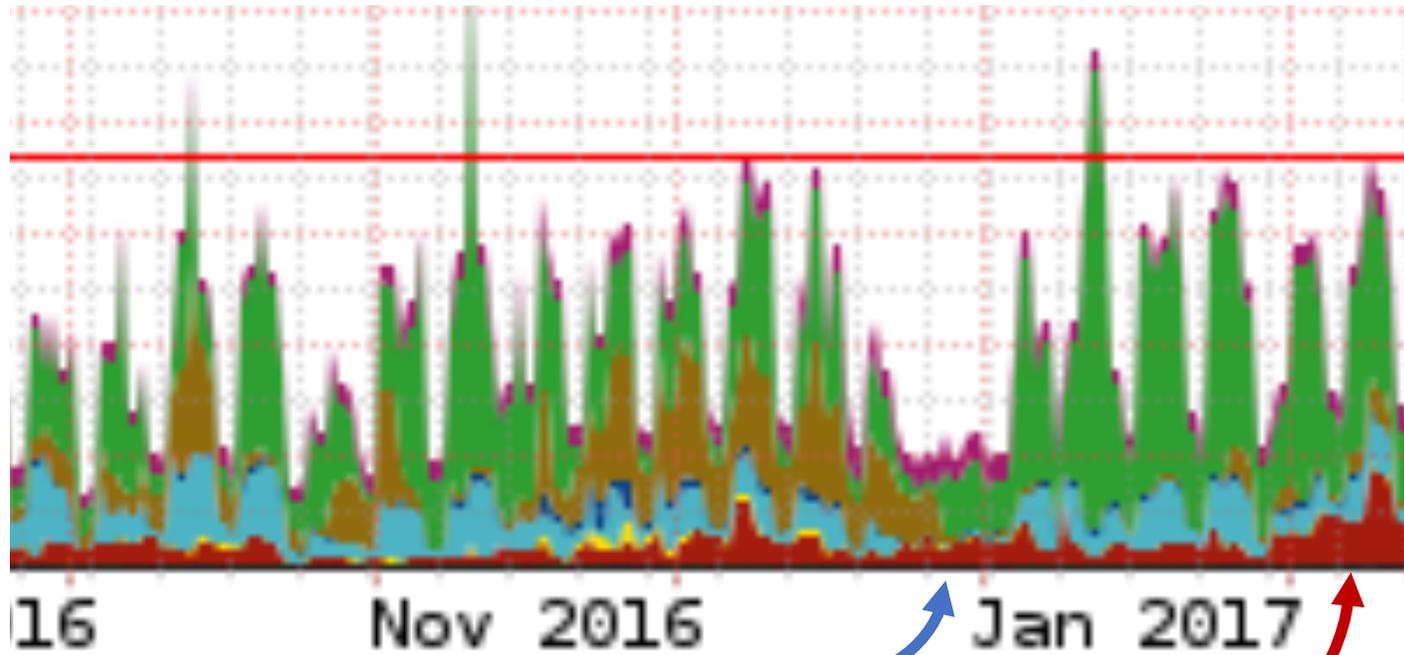
(sorry for the colours, not my favorite neither!)

AS20940 - Akamai International B.V. - IPV4



RRU700L / TOBI OETIKER

“Surprising” changes



Another change

Most likely as we do not peer in Amsterdam

Akamai most likely decided to consider our Leeds DNS to be “logically” in Manchester (due to low DNS query volume).

Not all Akamai clusters are born equal (content cached, size, etc.).

Akamai’s mapping is dynamic and changes to adapt to current conditions

- do not expect notifications
- the beast is “auto-magical”
- a bit like BGP ...

LINX is a great place to catch up with Akamai and discuss their current “policies”. Ronan could not be here today but feel free to contact him.

Conclusion

Installing more DNS resolvers ...

improved cache locality, making for a better end-user experience

- Moved some flows from London to Manchester
Reduced our London / Manchester core traffic
- Only worth it if you have multiple peering points
And eyeballs in different POPs
- If you want more traffic from regional exchanges
Setup some regional DNS
- DNS failure can cause traffic re-routing
Changing core link utilisation

