



IXLeeds Technical Update

IXLeeds 2 / 12 of September 2012

Thomas Mangin



Since IXLeeds 1

Work on our core services

Email

New switching platform

Looking Glass

Route Server

Behind the scene

Mailing Lists



LI

The initial Leeds Interconnect group

ops

for operational discussion

please make sure you are subscribed

members

for members discussion

very low volume

event-announce

new mailing list, subscribe if you want to know about IXLeeds' events.

<http://mailman.ixleeds.net/mailman/listinfo/members>

Contacting us by mail



noc@ixleeds.net , peering@ixleeds.net

any problems with the LAN or services

goes to andy and thomas

please do not send only one of us direct mails

support@ixleeds.net , info@ixleeds.net

any question about ixleeds

goes to the whole board (currently)

directors@ixleeds.net

the board email address

Migration to the Brocade RX8



Installation and Migration in March

Surprises during installation

- Back to front installation (wrong cold aisle)

- Front to side airflow intake for power

- Side to side air-flow intake for the switch

Overnight work

- No complication

- Port security works well (some of you know about it)

Thank you to

LINX for the donation and the configuration template / manuals

Andy and Craig for their help on the night

Right location



First Uplink



Completed
Installation



Switch Utilisation



Capacity

16x 10GB ports

96x 1GB ports

Utilisation

1x 100 Mb port

previous route collector (still doing IPv6)

21x 1 Gb port

12 1.5k VLAN members

2 9k VLAN members

7 Infrastructure (4x servers, 2x route servers, uplink)

2x 10 Gb member

new member port plugged (with PI)



Traffic




Traffic at the exchange

Still driven by a very few members

Most networks are eyeballs

Over 300 Mb this Autumn

This week around 100 Mb



Live graphs ?

changed a few time the backend

(witchtrial on previous server, cacti, going to FranceIX)

need to be merged to keep history

different RRD format .. not as simple as it sounds at first.

ETA: with SFlow or before

New Looking Glass



New tools available

<http://lg.ixleeds.net/>

by Daniel Piekacz (in the room) from Exa Networks
using ExaBGP (our current route-collector)

Publicly released last week

<https://code.google.com/p/gixlg/>

Already other public location using it ..

<http://www.freestone.net/lg/>

New Looking Glass



A few glitches to iron

Nothing major.

I am on it ...

Select: All None		Actions...		1 - 2 of 2		List	Grid
ID	Type	Status	Priority	Milestone	Owner	Summary + Labels	...
<input type="checkbox"/>	8 Defect	Accepted	Medium	----	thomas.mangin	› on sending sigterm, the child process terminate	
<input type="checkbox"/>	9 Defect	Accepted	Medium	----	thomas.mangin	› log rotation fails	

1 - 2 of 2 [CSV](#)

Still not completed.

IPv6 still on Cisco ..

Google “IPv6 vlan ubuntu bug” (now fixed)

Looking Glass



Public visibility of the collector

IXLeeds Looking Glass

router: request:
argument:

Node	Location	Country	RIR	AS name	ASN	Neighbor IP	IPv4/6	State	PfxRcd	Up/Down	Last update	Up since/last	Down since/last
sleek	Leeds	GB	RIPENCC	SLEEK-ASN	44009	91.217.231.11	4	up	7	5	21h28m47s	21h28m47s	2012-09-06 16:31:25
aaisp	Leeds	GB	RIPENCC	AS20712	20712	91.217.231.14	4	up	47	15	21h28m45s	21h28m46s	2012-09-06 17:28:25
fluidata	Leeds	GB	RIPENCC	FLUIDATA	39545	91.217.231.16	4	up	1	5	21h28m47s	21h28m47s	2012-09-06 16:31:25
janet	Leeds	GB	RIPENCC	JANET	786	91.217.231.17	4	up	257	5	21h28m43s	21h28m47s	2012-09-06 16:31:25
vaioni	Leeds	GB	RIPENCC	VAIONI	35575	91.217.231.18	4	up	19	5	21h28m45s	21h28m47s	2012-09-06 16:31:25
ask4	Leeds	GB	RIPENCC	ASK4	41230	91.217.231.19	4	up	21	5	21h28m47s	21h28m48s	2012-09-06 16:31:25

Looking Glass



Route information

IXLeeds Looking Glass

router: ixleeds collector

request: show ip bgp regexp [arg ASN] - kind of :)

argument: 30740 *

Execute

Node	Type	Location	Network	IPv4/6	Neighbor IP	AS path	Last seen
exa	peer	Leeds	46.226.168.0/21	4	91.234.18.12	30740 52200	2012-09-06 17:28:25
exa	peer	Leeds	192.175.48.0/24	4	91.234.18.12	30740 112	2012-09-06 17:28:31
Total number of prefixes 2							

[GIX looking glass](#)

Looking Glass



Route information

IXLeeds Looking Glass

router: ixleeds collector | request: show ip bgp [arg NETv4 or NETv6]

argument: 46.226.168.0 | Execute

Node	Location	Type	Network	IPv4/6	Neighbor IP	AS path	Last seen
exa	Leeds	peer	46.226.168.0/21	4	91.217.231.12	30740 52200	2012-09-08 16:49:40
exa	Leeds	peer	46.226.168.0/21	4	91.234.18.12	30740 52200	2012-09-08 16:49:40

Total number of prefixes 2

```
graph TD; A([IXLEEDS peering LAN]) --> B([AS30740 EXA-NETWORKS]); C([IXLEEDS 9k peering LAN]) --> B; B --> D([AS52200 M5-AS])
```

Looking Glass



Route information

IXLeeds Looking Glass

router request
ixleeds collector show ip bgp [arg NETv4 or NETv6]

argument
46.226.168.0 Execute

Node	Location	Type	Network	IPv4/6	Neighbor IP	AS path	Last seen
exa	Leeds	peer	46.226.168.0/21	4	91.217.231.12	30740 52200	2012-09-08 16:49:40
exa	Community: 30740:65104 Extended community: Origin: egp Nexthop: 91.217.231.12						16:49:40

Total number of routes: 1

```
graph TD; A([AS30740 EXA-NETWORKS]) --> B([AS52200 M5-AS]);
```

GIXLG in a network



Exa Looking Glass

router: request: argument:

Node	Location	Country	RIR	AS name	ASN	Neighbor IP	IPv4/6	State	PfxRcd	Up/Down	Last update	Up since/last	Down since/last
Juniper MX80-1	Manchester	GB	RIPENCC	EXA-NETWORKS	30740	82.219.0.68	4	up	420329	1	1s	4h56m54s	never
Juniper MX80-2	Manchester	GB	RIPENCC	EXA-NETWORKS	30740	82.219.0.69	4	up	420329	1	1s	4h56m54s	never
Juniper MX80	London	GB	RIPENCC	EXA-NETWORKS	30740	82.219.0.71	4	up	420329	1	1s	4h56m54s	never
Cisco 7204VXR	Cottingley	GB	RIPENCC	EXA-NETWORKS	30740	82.219.0.76	4	up	18	1	4h56m52s	4h56m54s	never
Cisco 7301	Manchester	GB	RIPENCC	EXA-NETWORKS	30740	82.219.0.77	4	up	534	1	40s	4h56m54s	never
Cisco 7201	London	GB	RIPENCC	EXA-NETWORKS	30740	82.219.0.78	4	up	360	1	38s	4h56m54s	never
Juniper EX4200	Leeds	GB	RIPENCC	EXA-NETWORKS	30740	82.219.0.79	4	up	163	1	4h56m51s	4h56m54s	never
Juniper MX80	Manchester	GB	RIPENCC	EXA-NETWORKS	30740	82.219.0.88	4	down	0	0	never	never	never
Juniper MX80	London	GB	RIPENCC	EXA-NETWORKS	30740	82.219.0.90	4	down	0	0	never	never	never
Juniper MX80	Leeds	GB	RIPENCC	EXA-NETWORKS	30740	82.219.0.92	4	up	420330	1	1s	4h50m17s	never
Juniper MX80	Manchester	GB	RIPENCC	EXA-NETWORKS	30740	2a02:b80:18:0:b3	6	down	0	0	never	never	never
Juniper MX80	London	GB	RIPENCC	EXA-NETWORKS	30740	2a02:b80:1a:0:b3	6	down	0	0	never	never	never
Juniper MX80	Leeds	GB	RIPENCC	EXA-NETWORKS	30740	2a02:b80:1c:0:b3	6	up	10256	1	1s	4h50m17s	never
Juniper MX80-1	Manchester	GB	RIPENCC	EXA-NETWORKS	30740	2a02:b80:4:0:b3	6	up	10256	1	1s	4h40m29s	never
Juniper MX80-2	Manchester	GB	RIPENCC	EXA-NETWORKS	30740	2a02:b80:5:0:b3	6	up	10255	1	1s	4h41m47s	never
Juniper MX80	London	GB	RIPENCC	EXA-NETWORKS	30740	2a02:b80:7:0:b3	6	up	10255	1	1s	4h44m7s	never

Total number of neighbors 16

Visualise your Paths



Exa Looking Glass

router request: ExaRC show ip bgp [arg NETv4 or NETv6]

argument: 8.8.8.8

Node	Location	Type	Network	IPv4/6	Neighbor IP	Next hop	Dest. node	Dest. type	Dest. loc.	AS path	Last seen
Juniper MX80-1	Manchester	core	8.8.8.0/24	4	82.219.0.68	82.219.0.71	Juniper MX80	core	London	30740 15169	2012-09-07 10:41:22
Juniper MX80-2	Manchester	core	8.8.8.0/24	4	82.219.0.69	82.219.0.71	Juniper MX80	core	London	30740 15169	2012-09-07 10:41:57
Juniper MX80	London	core	8.8.8.0/24	4	82.219.0.71	195.66.224.125	LINX Juniper	peering LAN	London	30740 15169	2012-09-07 10:41:52
Juniper MX80	Leeds	core	8.8.8.0/24	4	82.219.0.92	82.219.0.71	Juniper MX80	core	London	30740 15169	2012-09-07 10:48:28
Juniper MX80-1	Manchester	core	8.0.0.0/8	4	82.219.0.68	194.222.89.0	KPN	transit	Manchester	30740 286 3356	2012-09-07 10:42:35
Juniper MX80-1	Manchester	core	8.0.0.0/9	4	82.219.0.68	194.222.89.0	KPN	transit	Manchester	30740 286 3356	2012-09-07 10:42:35
Juniper MX80-2	Manchester	core	8.0.0.0/8	4	82.219.0.69	82.219.0.68	Juniper MX80-1	core	Manchester	30740 286 3356	2012-09-07 10:41:37
Juniper MX80-2	Manchester	core	8.0.0.0/9	4	82.219.0.69	82.219.0.68	Juniper MX80-1	core	Manchester	30740 286 3356	2012-09-07 10:41:37
Juniper MX80	London	core	8.0.0.0/8	4	82.219.0.71	82.219.0.68	Juniper MX80-1	core	Manchester	30740 286 3356	2012-09-07 10:41:15
Juniper MX80	London	core	8.0.0.0/9	4	82.219.0.71	82.219.0.68	Juniper MX80-1	core	Manchester	30740 286 3356	2012-09-07 10:41:15
Juniper MX80	Leeds	core	8.0.0.0/8	4	82.219.0.92	82.219.0.68	Juniper MX80-1	core	Manchester	30740 286 3356	2012-09-07 10:48:29
Juniper MX80	Leeds	core	8.0.0.0/9	4	82.219.0.92	82.219.0.68	Juniper MX80-1	core	Manchester	30740 286 3356	2012-09-07 10:48:29

Total number of prefixes 12

```

graph TD
    JML[Juniper MX80 core Leeds] --> JMM[Juniper MX80-1 core Manchester]
    JML --> JMLC[Juniper MX80 core London]
    JMM --> JML
    JMM --> JMLC
    JMM --> JMM2[Juniper MX80-2 core Manchester]
    JMM2 --> JMM
    JMM2 --> JMLC
    JMLC --> JMM
    JMLC --> JMM2
    JMLC --> KPM[KPN transit Manchester]
    KPM --> AS286[AS286 KPN]
    AS286 --> AS3356[AS3356 LEVEL3]
    JMLC --> JMLC2[Juniper MX80 core London]
    JMLC2 --> LINX[LINX Juniper peering LAN London]
    LINX --> AS15169[AS15169 GOOGLE]
  
```


More Specific 8.8.8.0/24

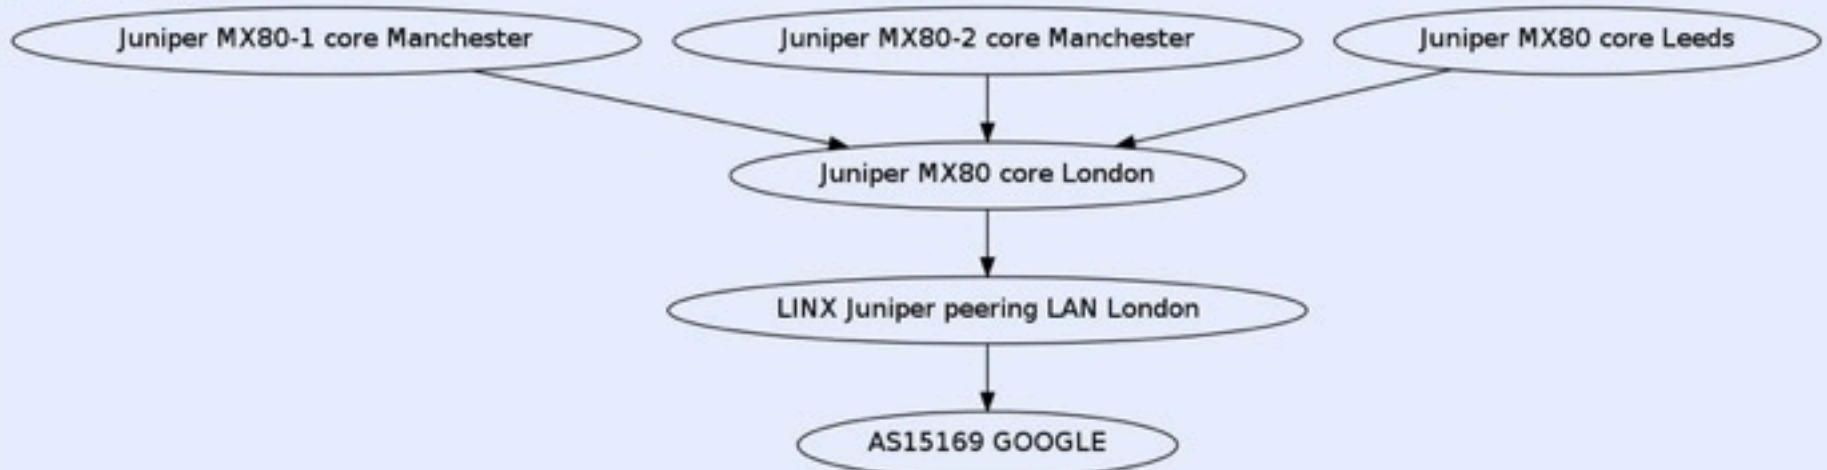


Exa Looking Glass

router request
ExaRC show ip bgp [arg NETv4 or NETv6]
argument
8.8.8.0/24 Execute

Node	Location	Type	Network	IPv4/6	Neighbor IP	Next hop	Dest. node	Dest. type	Dest. loc.	AS path	Last seen
Juniper MX80-1	Manchester	core	8.8.8.0/24	4	82.219.0.68	82.219.0.71	Juniper MX80	core	London	30740 15169	2012-09-07 10:41:22
Juniper MX80-2	Manchester	core	8.8.8.0/24	4	82.219.0.69	82.219.0.71	Juniper MX80	core	London	30740 15169	2012-09-07 10:41:57
Juniper MX80	London	core	8.8.8.0/24	4	82.219.0.71	195.66.224.125	LINX Juniper	peering LAN	London	30740 15169	2012-09-07 10:41:52
Juniper MX80	Leeds	core	8.8.8.0/24	4	82.219.0.92	82.219.0.71	Juniper MX80	core	London	30740 15169	2012-09-07 10:48:28

Total number of prefixes 4



Route Servers



Finally ... They are ready

Thank you to

 **LONAP** for sharing their work with us

Andy for installing the server and configuring them ...

You can connect !

Collector : 91.217.231.1 / AS 51526

RS1 : 91.217.231.2 / AS 57932

RS2 : 91.217.231.3 / AS 57932

Not yet available on the 9k LAN

“Preempting : No Martin, no IPv6. At least not yet”

Behind the scene Intranet



Member	<input type="text" value="Exa Networks"/>
Connection	<input type="text" value="rx1.dc2.aql"/> <input type="text" value="5"/> <input type="text" value="12"/> <input type="text" value="1 Gb"/>
Address	<input type="text" value="27-29 mill field road
cottingley business park"/>
City	<input type="text" value="Bingley"/>
County	<input type="text" value="West Yorkshire"/>
Postcode	<input type="text" value="BD16 1PY"/>
Country	<input type="text" value="United Kingdom"/>
Peering	<input type="text" value="Exa Peering"/>
Technical	<input type="text" value="Exa NOC"/>
Administrative	<input type="text" value="Thomas Mangin"/>
Financial	<input type="text" value="Exa Accounts"/>
Commercial	<input type="text" value="Thomas Mangin"/>
Technical	<input type="text" value="Thomas Mangin"/>
Back	

Behind the scene Intranet



We are missing contact information for a few members, we will be contacting you shortly.

Exa Peering		role contact
Phone	<input type="text" value="1"/>	<input type="text" value="00 44 845 145 1234"/>
Email	<input type="text" value="1"/>	<input type="text" value="peering@exa-networks.co.uk"/>
Back		

Behind the scene
continued ...



Quickly* to not be too boring

Paper .. work

- The connection policy ratified
- RIPE resources request and transfer
- Documentation (network, servers, ...)

Server .. work

- New servers installation and maintenance
 - WEB, DNS, NTP, SYSLOG, MAIL/Mailing Lists, ...
- Offsite backups

Website .. work

- Website (DB integration) whizz-kid anywhere ?

Late .. work

- quite a _lot_ on our todo lists



* Yes I can speak even faster



time is limited .. problem happens

SFLOW status


LAN issues

and lesson learned

SFLOW



A work in progress.

Thank you very much to  FranceIX

For sharing their work with us ..

is there any IX we did not borrow something from ?

Planning to make their work generic.

And release it to every Euro-IX members

Some work left

The code is good, need cleanup and making generic

Issues



One global LAN issue

- all the port went down
- required a reboot of the switch
- found an issue with remote syslog ...

Many port security violation

A few member suffering from MAC leaks

```
interface ethernet 5/12
port-name "PEER:: Exa Networks"
no fdp enable
no flow-control
sflow forwarding
port security
maximum 6
restrict-max-deny 6
secure-mac-address 6487.885c.fd74 402
secure-mac-address 6487.885c.fd74 401
```

Route Collector

- Two separate issues (one in April) not service affecting (one last week)
- Fixed quickly (but some members left their session down)
- left members without session while it was updated / fixed



What did we learn

Need for better (external) monitoring

Members were faster to detect the fault than we were

Better out of band solution

Make sure we have console OOB

The right console cables is now left on-site

We could not access the switch before the reboot :(



New services

9K VLAN

Private Interconnect

More services ?



9,000 Bytes MTU VLAN

Inspiration : [draft-mlevy-ixp-jumboframes-00](#)

Took some convincing to get RIPE to agree ..

Allows to pass encapsulated traffic unfragmented

L2TP, SAN, ...

Requires a new port

Allocated IP ranges are:

Protocol	MTU Size	IP Address Block
IPv4	1,500 Bytes	91.217.231.0/24
IPv6	1,500 Bytes	2001:7F8:67:0::/64
IPv4	9,000 Bytes	91.234.18.0/24
IPv6	9,000 Bytes	2001:7F8:67:9::/64

No extra cost - Just ask for it

New Service



Member to Member PI

Provided by many IXes

https://www.euro-ix.net/tools/ixp_matrix

“vlan services”

Possible on

existing port as a VLAN

most practical, preferred route

on a dedicated port

for free (as long as we have no port pressure constraint)

Like normal peering LAN, two options

IPv4 / IPv6

1500 / 9000 Bytes MTU

Moving Forward ..



Starting new services ?

Private User Group

Point to Multi-Point

Still BGP between different speakers

Provide more value to our members

starting partnership to become as well a platform for services

A marketplace for transit ?

Allowing different joining condition for foreign networks ?

Require a consultation

Board strategic review on new services

Presentation to members for a vote



Good of the internet

L-ROOT

M-LAB

News Service

More ?

Looking Forward



IX helping to keep traffic local

one of the reason for peering

IX helping with resilience

peering is distributed

IX helping internet research

helps our members

For the good of the internet



Providing connectivity to friendly organisation

L-ROOT DNS Server

Connecting an AnyCast DNS server

Provide resilience for DNS

Google “godaddy outage million sites”

M-LAB

Measurement Lab (M-Lab) is an open, distributed server platform for researchers to deploy Internet measurement tools.

For the good of the t'interweb



News service

requested by a member

access to usenet news to your customers via IXLeeds

binary free to begin with

once online, point your DNS to it :D

Any more ?

mirror of popular sites ? Distribution, etc.

others

Questions?
thomas@ixleeds.net

