

Exporting Flow Telemetry For Fun and Profit

October 2016
RIPE 73, Madrid



i n e x
i n t e r n e t n e u t r a l e x c h a n g e

Nick Hilliard

CTO

nick@inex.ie



Sflow Telemetry - Problem Statement

i n t e r n e t n e u t r a l e x c h a n g e

DDoS is causing problems on the Internet and hitting international media

When DDoS traffic passes over an IXP, it's not traceable to the source ASN

IXPs have no visibility into what constitutes DDoS/non-DDoS traffic

Almost no high-end routers export MAC address information in netflow

Even if the IXP has sflow telemetry from IXP infrastructure, it cannot be exported to IXP participants



Sflow Telemetry - Problem Statement

i n e x
i n t e r n e t n e u t r a l e x c h a n g e

Sflow export from IXP to participants is not possible because:

- sflow is implemented in a container format with multiple records per packet
- no software which splits up sflow packets and filters on individual records
- it's not ok to export sflow records to participant A for data between participants B and C

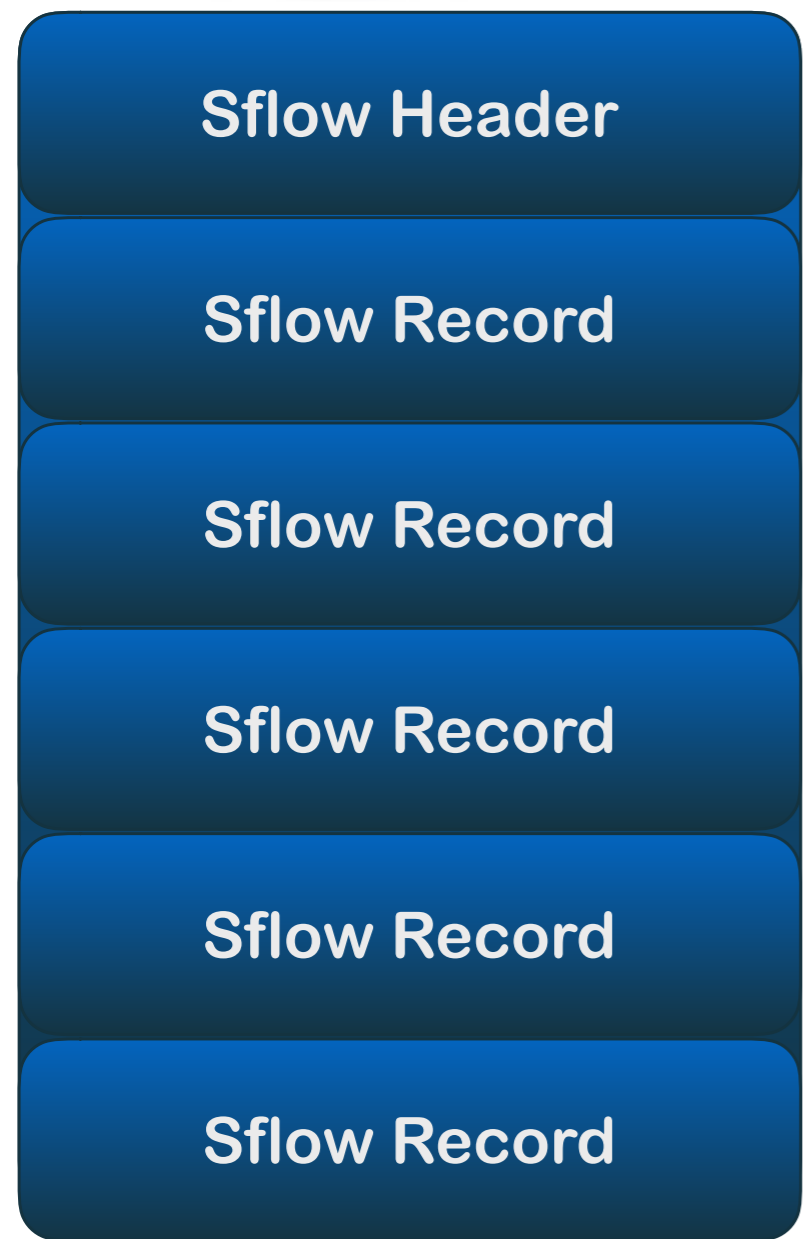


Sflow Telemetry - Packet Structure






Sflow Multiplexing



B
C
B
C
C

Packet X



Packet Y



src mac B
src mac B
src mac C
src mac C
src mac C



i n e x
i n t e r n e t n e u t r a l e x c h a n g e

PMacct

INEX talked to Paolo Lucente, who wrote PMacct

No support for this feature, but it had a framework which made it viable

Paolo wrote some code and sflow demuxing + filtering is now supported in pmacct



i n t e r n e t n e u t r a l e x c h a n g e

PMacct Configuration

```
sfacctd_ip: 193.242.111.xxx
sfacctd_port: 6343
!
plugins: tee[sflow_receiver]
!
pre_tag_map[sflow_receiver]: /srv/ixpm/pretag.map
maps_entries: 3000
maps_index: true
!
tee_receivers[sflow_receiver]: /srv/ixpm/receivers.lst
tee_max_receiver_pools[sflow_receiver]: 1000
!
tee_dissect_send_full_pkt[sflow_receiver]: true
tee_transparent[sflow_receiver]: true
```



PMacct Configuration

/srv/ixpm/pretag.map

```
set_tag=32          src_mac=02:1d:b5:c3:e8:2a  
set_tag=32          dst_mac=02:1d:b5:c3:e8:2a
```

/srv/ixpm/receivers.lst

```
id=32              ip=192.168.237.5:6343          tag=32
```


Search...



IXP CUSTOMER ACTIONS

Customers

Interfaces

Users

Contacts

Colocated Equipment

Meetings

IXP ADMIN ACTIONS

Infrastructures

Locations

Cabinets

Switches

IP Addressing

MAC Addresses

Vendors

Console Server

Home / Sflow Receivers / Add New Sflow Receiver



Destination IP

192.168.237.5

Destination Port

6343

Add

Cancel

[Save Changes](#)[Return to Customer Overview](#)[Advanced Options](#)

IXP ADMIN

ACTIONS

[Infrastructures](#)[Locations](#)[Cabinets](#)[Switches](#)[IP Addressing](#)[MAC Addresses](#)[Vendors](#)[Console Server](#)[Connections](#)[VLANs](#)[IRRDB](#)[Configuration](#)[Route Server](#)[Prefixes](#)

IXP STATISTICS

[Member](#)[Statistics -](#)[Graphs](#)[Member](#)[Statistics - List](#)[Member Logos](#)[League Table](#)



Physical Interfaces +

Location	Peering Port	Fanout Port	Speed/Duplex	
Equinix Kilcarbery	swi1-deg1-3::1:44		1000/full	 

VLAN Interfaces +

VLAN Name	VLAN ID	IPv4 Address	IPv6 Address	
Peering VLAN #1	10	193.242.111.6	2001:7f8:18::6	 

Sflow Receivers +

Target IP	Target Port	
192.168.237.5	6343	 



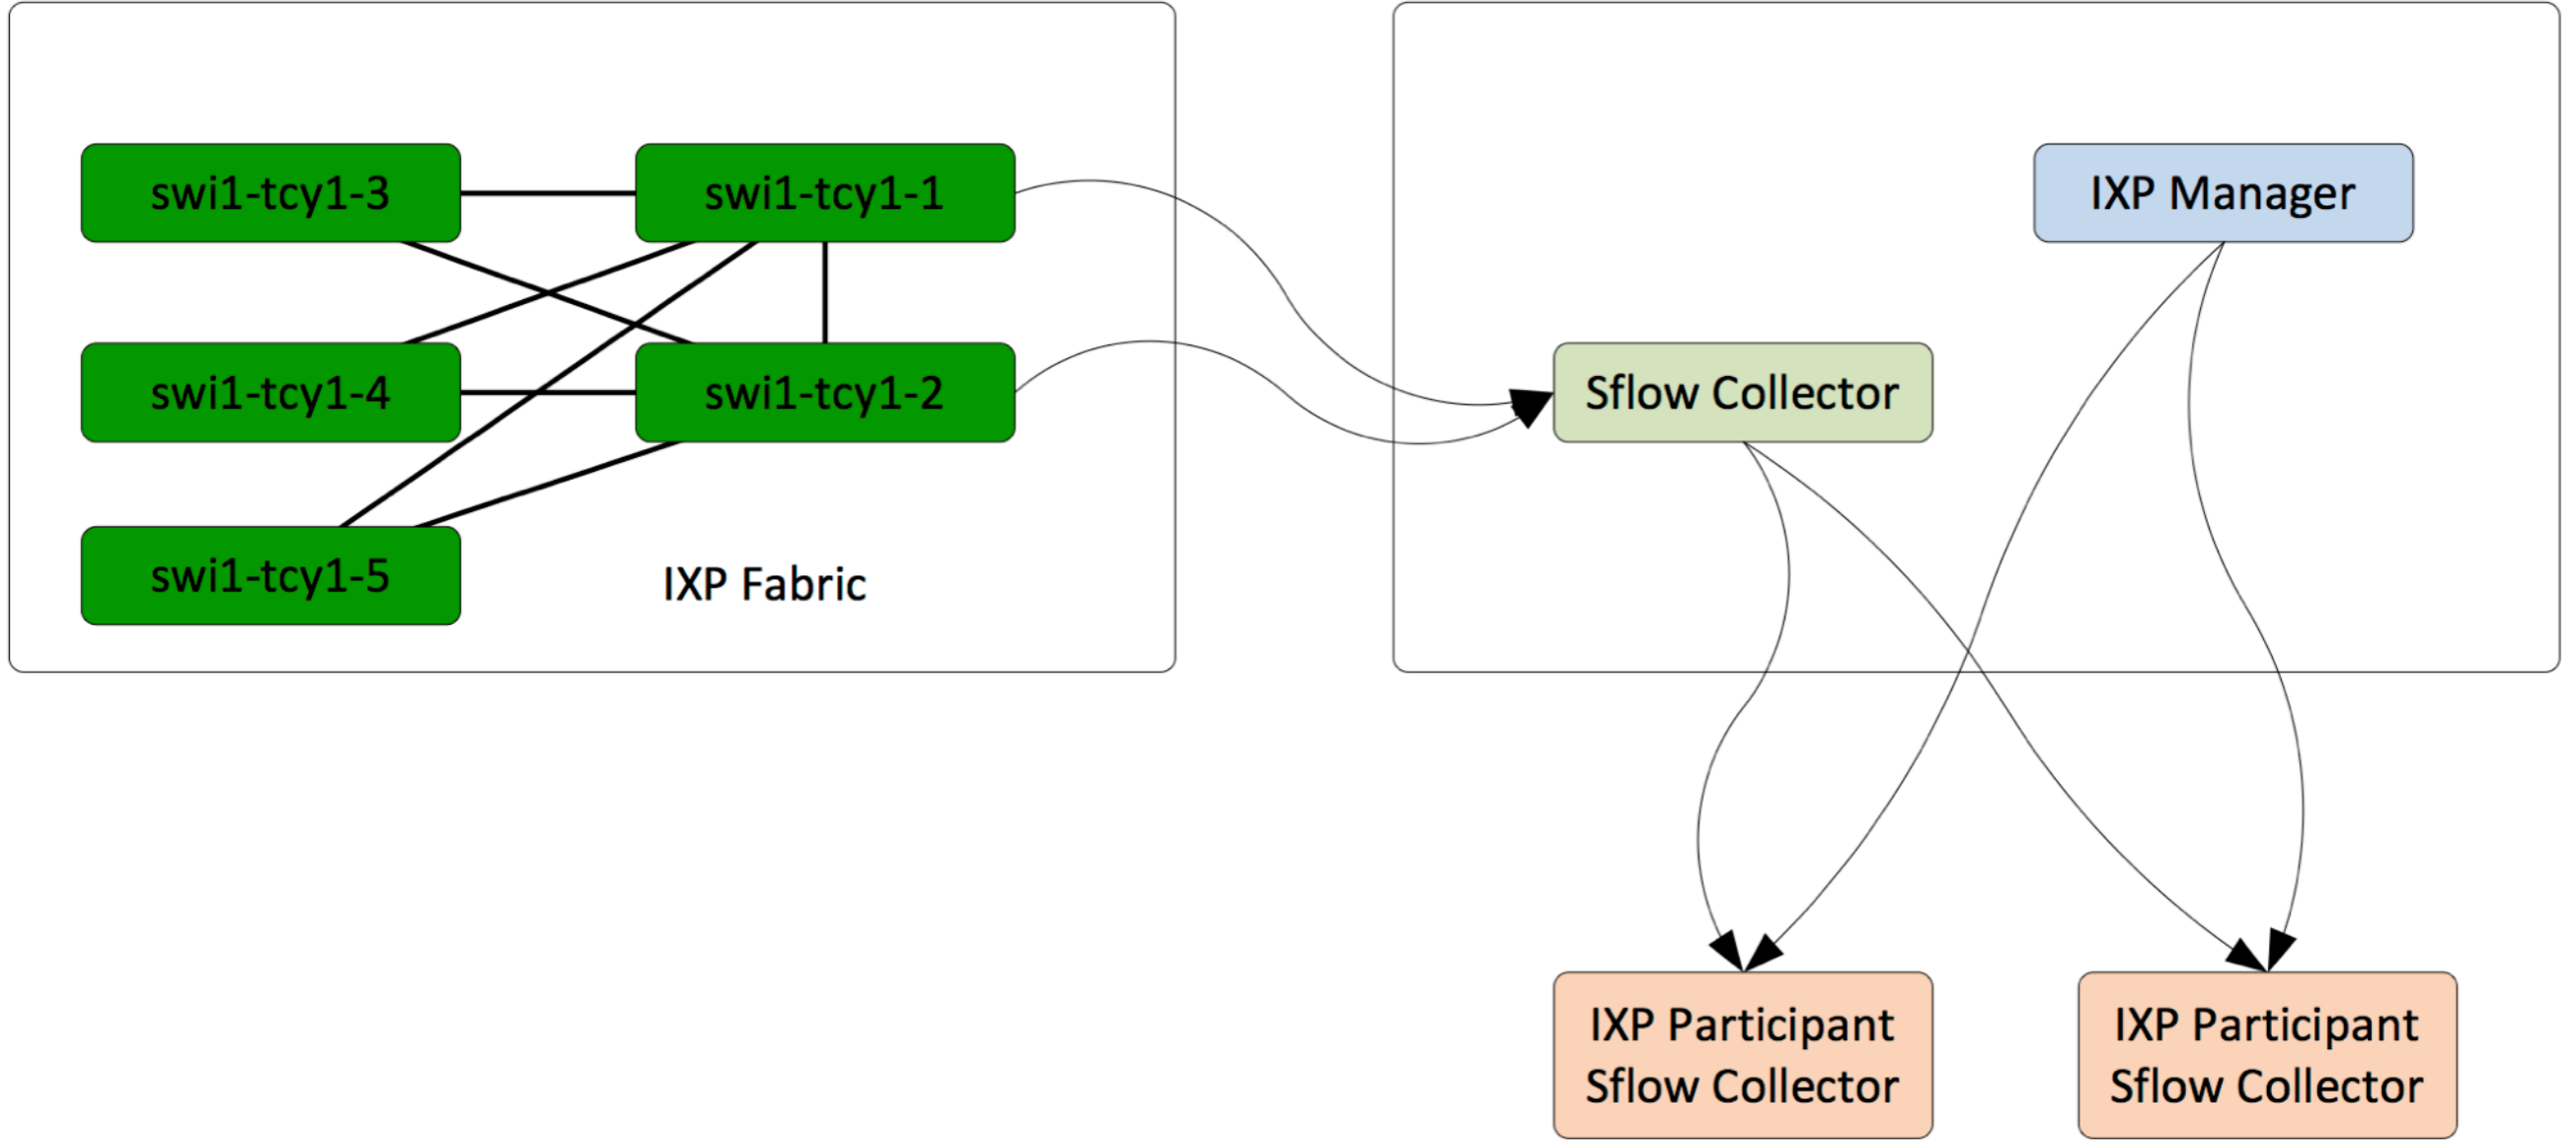
i n e x
i n t e r n e t n e u t r a l e x c h a n g e

Hooking it all Together

Euro-IX JSON export schema supports export of MAC addresses

IXP participants need to run sflow collectors

... but most netflow collectors also support sflow





Result

Allows IXP participants to perform analysis of their traffic flows at IXP

No privacy issues

Can be used to reliably trace the source of spoofed flows

Live data, suitable for input into reactive traffic management

Currently in pilot phase at INEX

Both PMacct and IXP Manager code available on github

Sflow Data Export from IXPs



i n e x
i n t e r n e t n e u t r a l e x c h a n g e

That's all folks...

...any questions?