



# sFlow implementation at AMS-IX

NANOG 40  
Elisa Jasinska  
[elisa.jasinska@ams-ix.net](mailto:elisa.jasinska@ams-ix.net)

# Agenda

- What is sFlow?
- AMS-IX requirements
- Hardware specifics
- Software
- Results and usage
- Future plans

# What is sFlow?

- Capture traffic data in switched or routed networks
- Sampling technology
- Datagram format standard defined in RFC 3176
- Implemented on a wide range of devices (Foundry, Force10, Extreme...)

# What is sFlow?

- Not everything is sampled information
- Two different types provided by the datagram format:
  - Flow samples
  - Counter samples

# What is sFlow?


- Flow samples
  - Whole captured packet (L2-L7)
  - Defined sampling rate (eg. one out of 8192)
- Counter samples
  - Interface counters (octets/packets/errors/...)
  - Polling interval (eg. 30 sec.)

# AMS-IX Requirements

- Use flow samples to show member to member traffic statistics
  - Operates only on layer 2
  - One MAC address per member
- Show other information, eg. ether type
- Use counter samples to show interface statistics



# AMS-IX Requirements

- High performance demands
- 260 Gbps - 40 Mpps
- Sampling rate 8192 → ca. 4800 samples per second



# AMS-IX Requirements


- Issues with MRTG



- Spikes due to CPU load





# Hardware Specifics




# Hardware Specifics

- Foundry - inbound traffic
  - Packets dropped by the switch still counted
- Force10 - outbound sampling




# Hardware Specifics

- Counter samples with fixed polling interval
- Different and not configurable arrival times



# Hardware Specifics

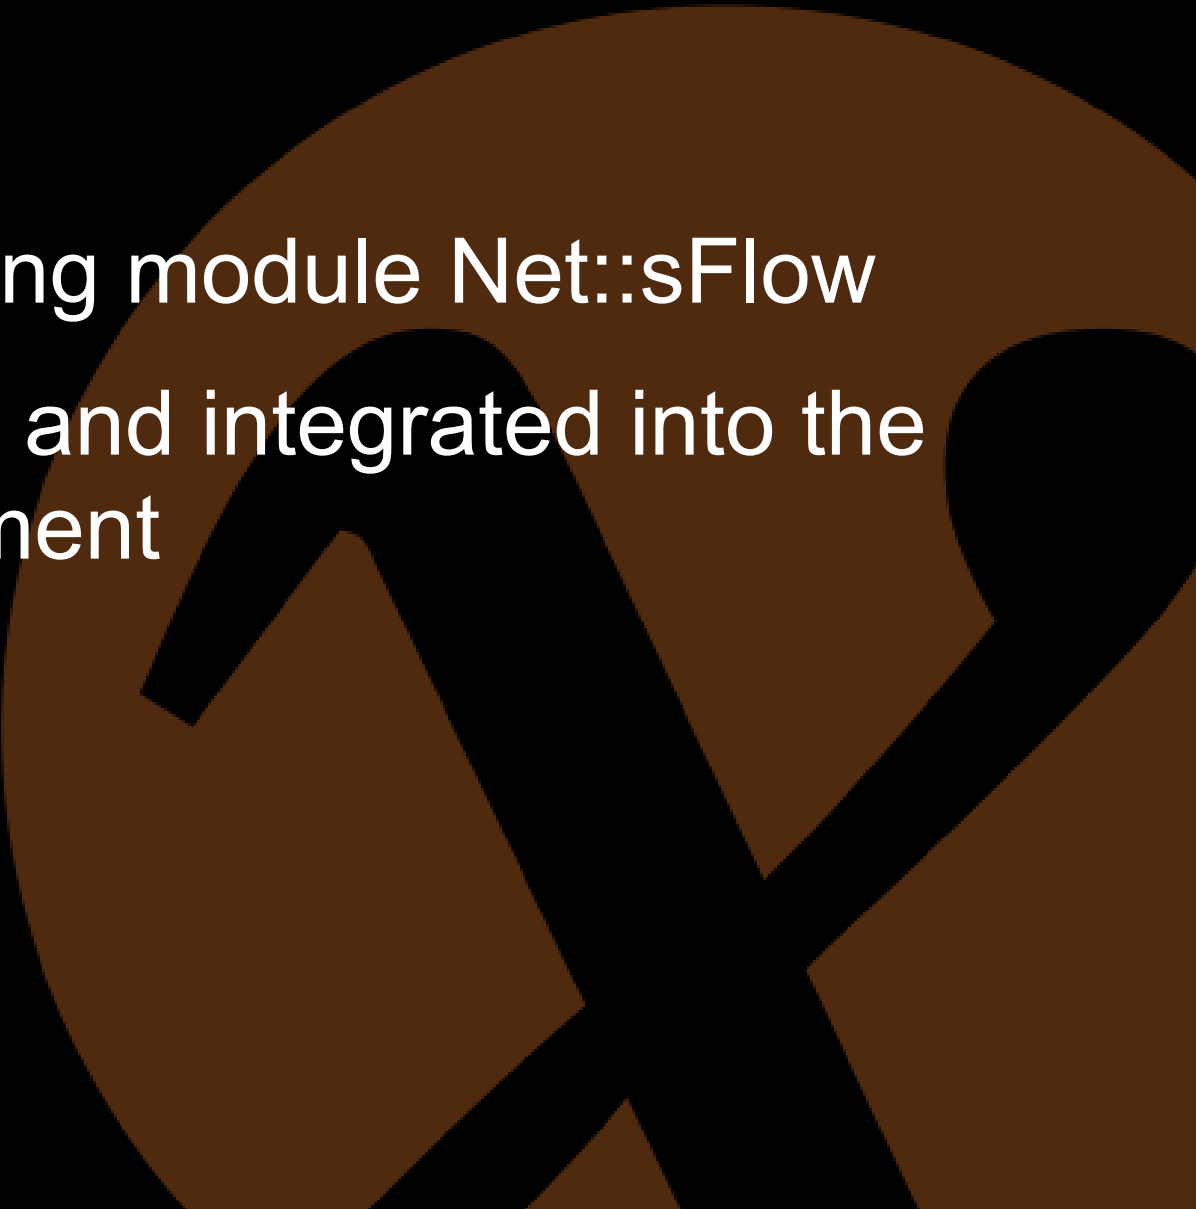
- To accumulate traffic correctly high interval needed



# Software

- InMon – sflowtool
- Pmacct
- InMon – Traffic Sentinel
- libsflow / sflowd
- ...

# AMS-IX Software

- Written in PERL
  - Based on decoding module Net::sFlow
  - Fully customized and integrated into the AMS-IX environment
- 


# AMS-IX Software

- CPU usage growing linearly with amount of packets/samples
- I/O performance feasible
  - Preprocessing the data
  - Only storing needed information
  - Currently writing 50 000 files


# Results and Usage

- Accuracy

## sFlow



## SNMP





# Results and Usage

## Fnord Internet B.V.


Switchport: 23@switch01  
 IP: 192.168.45.131  
 AS: 25538  
 Route server: yes

Sorted by bps Top 10 Sorted by "Total" (both directions) Go

### Note:

The graphs are sorted by the sum of bps or pps over the last 24 hours, a single peak will not necessarily make the peer appear on top of the list.

1.




### Foobar Industries INC

IP: 192.168.45.134  
 AS: 250  
 Switch: switch02  
 Route server: no

AS details:  
[AS 25538 to AS 250](#)

2.

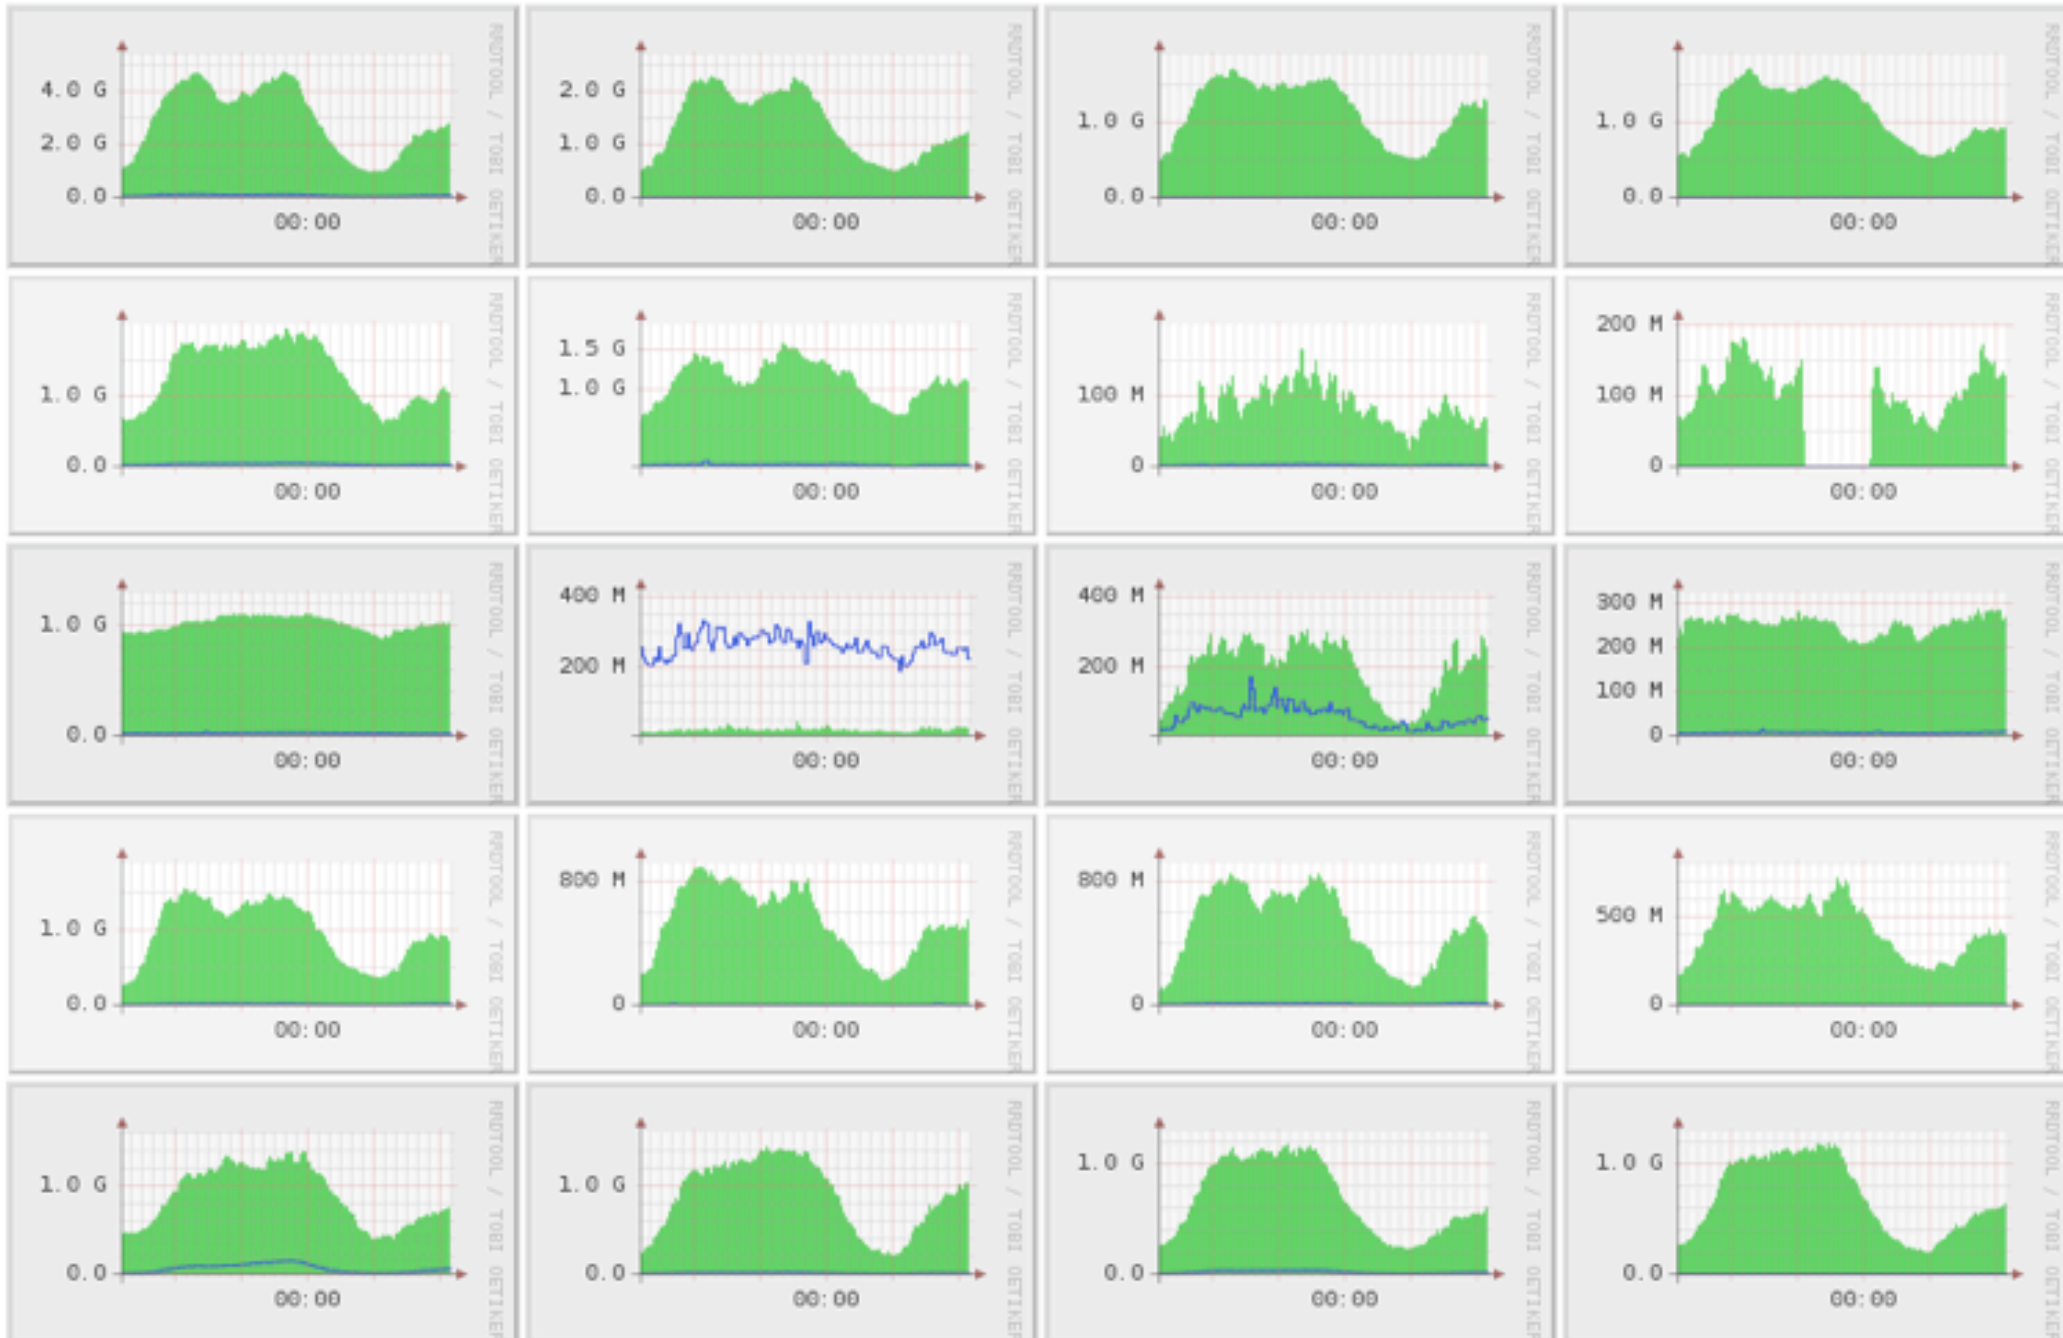


### Network GmbH



IP: 192.168.45.102  
 AS: 248  
 Switch: switch03  
 Route server: yes

AS details:  
[AS 25538 to AS 248](#)

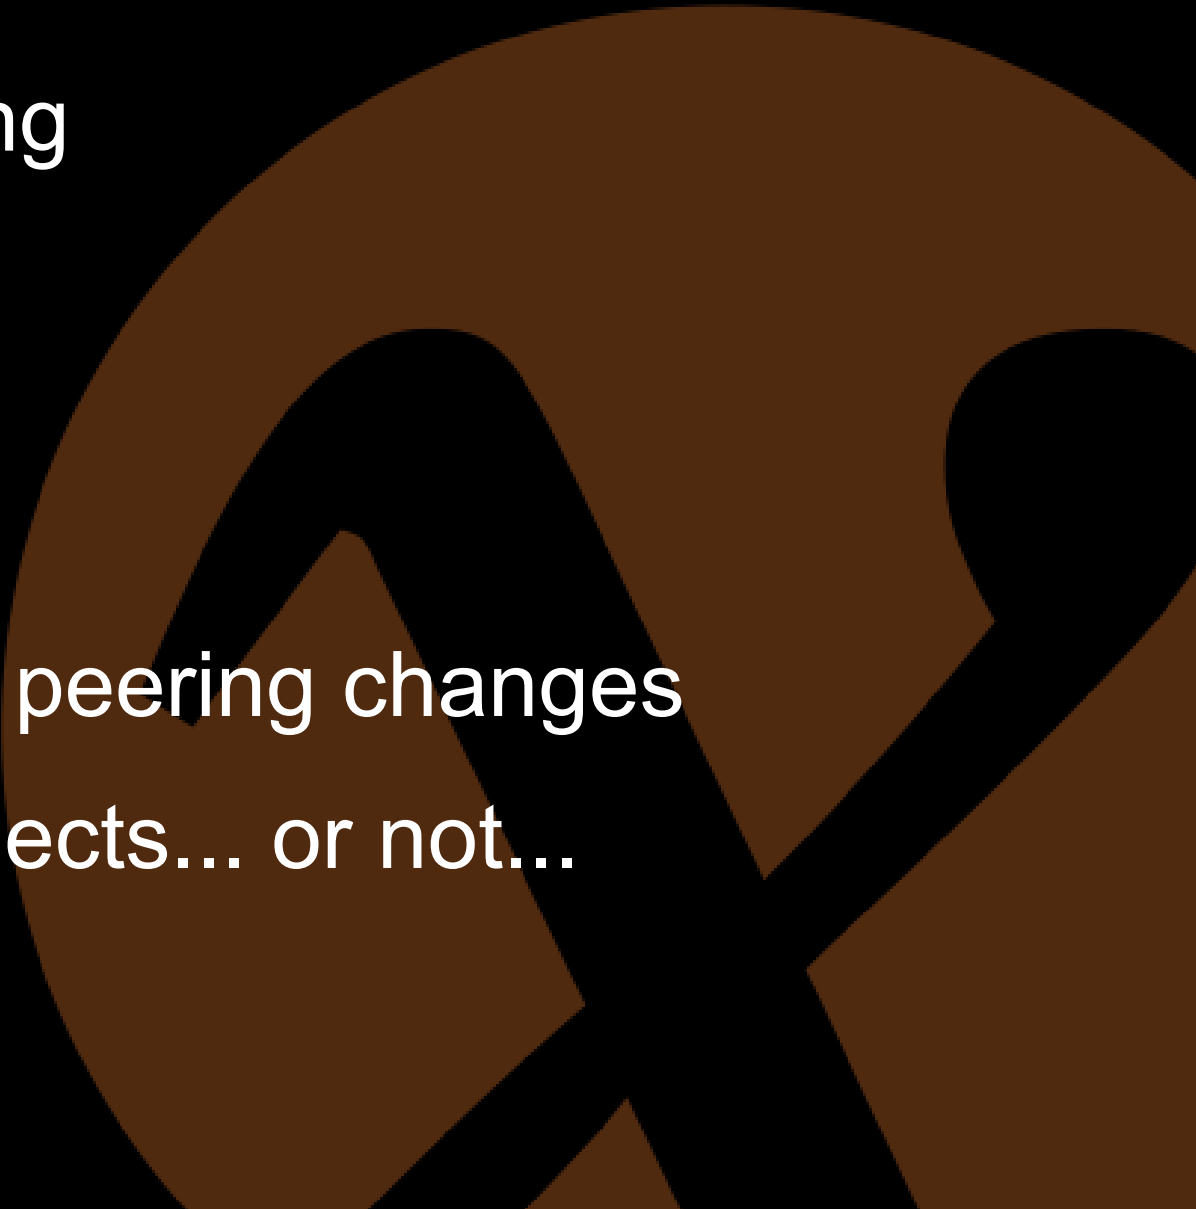
# sFlow Peering Matrix



# Results and Usage



# Results and usage

- Traffic engineering
    - Members
    - AMS-IX NOC
  - Debugging
  - Detailed view on peering changes
  - Private interconnects... or not...
- 

# Future plans

- Automated detection of...
  - Peerings
  - Outages
  - Traffic shifts
  - Fully utilized links
  - ...

# Thanks for listening!

## Questions?

[elisa.jasinska@ams-ix.net](mailto:elisa.jasinska@ams-ix.net)