sFlow

Elisa Jasinska
elisa.jasinska@ams-ix.net

Agenda

- What is sFlow?
- AMS-IX requirements
- Existing software solutions
- Performance issues
- Software used at AMS-IX
- Privacy

What is sFlow?

- Monitoring networks
- Cisco IOS - NetFlow
- Sampling mechanism, not “touching” every packet
- Applicable to high speed networks (>= 1GE)

What is sFlow?

- sFlow datagrams sent via UDP
- Datagram format standard defined in RFC 3176
- Implemented on a wide range of devices (Foundry, Force10, Extreme...)

What is sFlow?

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

What is sFlow?

- What do you need?
  - Hardware supporting sFlow
  - Central server to collect the data
  - Software to analyze the received data

What is sFlow?

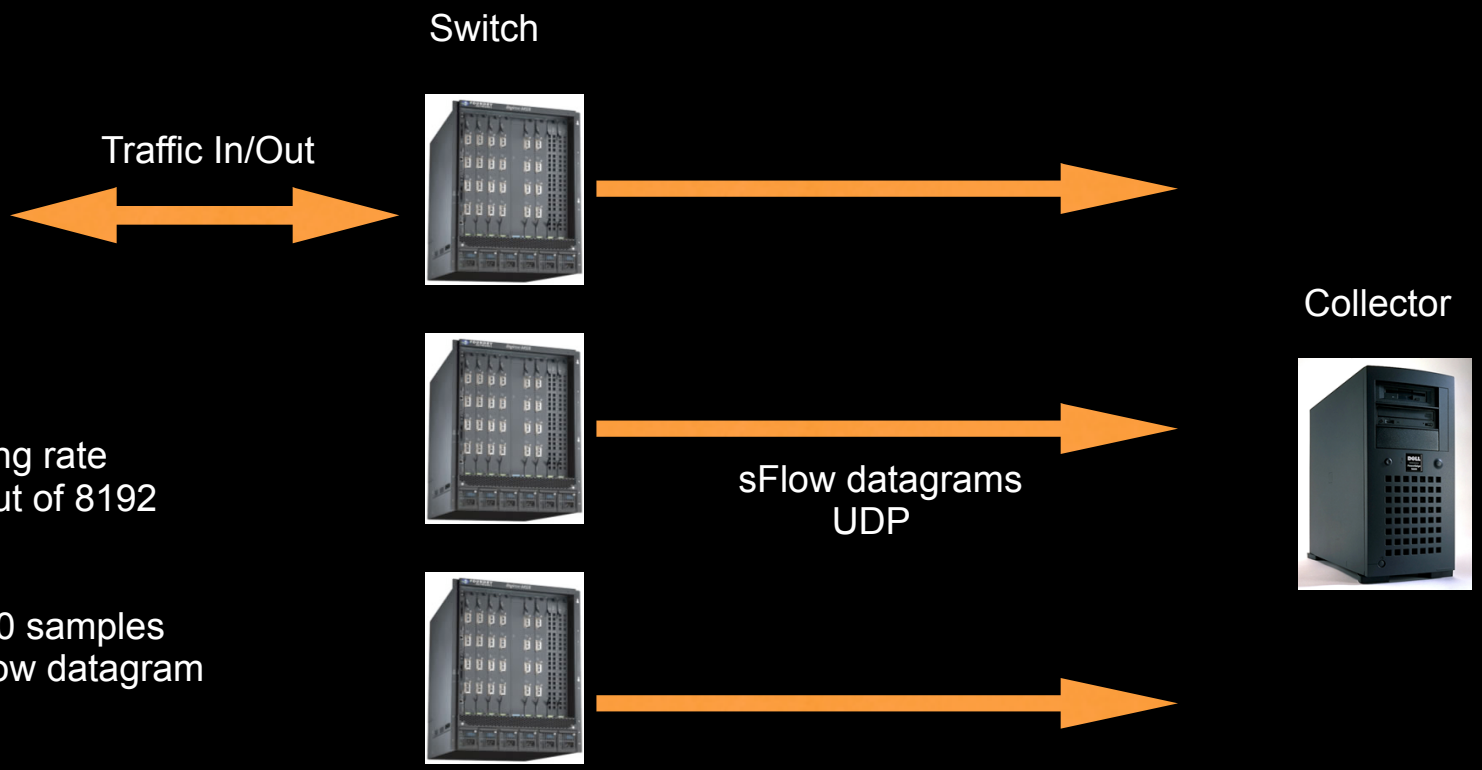

AMS-IX requirements

- Use flow samples to:
  - Provide member-to-member traffic information
  - See growth (or lack) of of IPv6
- Due to high throughput a very efficient system is required

Existing software solutions

- Free software:
  - InMon – sflowtool
  - Pmacct
  - sFlow2MySQL
- Commercial:
  - InMon – Traffic Sentinel

Existing software solutions

- Issues with existing software
  - Saves each sample to DB
  - No caching or preprocessing possible
  - Graphing with RRDtool
    - overhead due to data export to RRD
    - same data saved twice

Performance issues

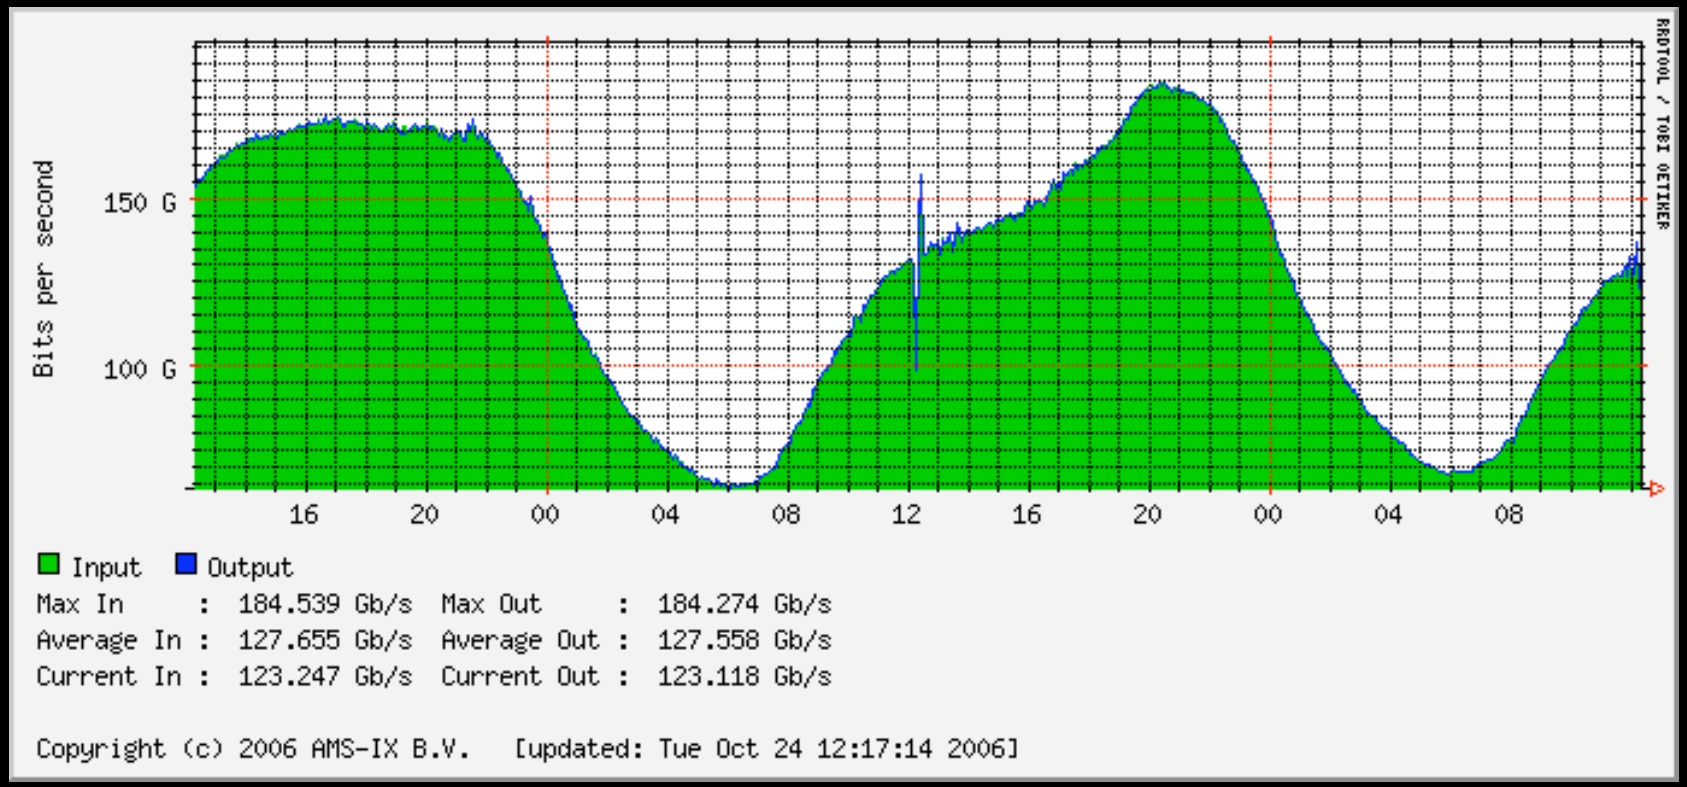

Performance issues

- Traffic up to 180 Gb/s (30 Mpps)
- ca. 3500 samples per second
- Cannot store each sample in a DB

Software used at AMS-IX

- Written in PERL
  - Easy to understand
  - Good integration with RRDtool
  - Due to PERL's re-use architecture (modules) lots of subtasks have already been programmed
  - Largest common denominator of a language understood at the AMS-IX NOC

Net::sFlow

- Decodes sFlow datagrams
- Supports sFlow version 2/4 and 5
- Single (exportable) function, decode()
- Available on CPAN

sFlow daemon

- Based on module Net::sFlow
- Receives UDP datagrams
- Analyzes the information
- Stores data to RRD files

Performance Results

- CPU usage while decoding sFlow datagrams
- Growing linearly with amount of packets / samples

Performance Results

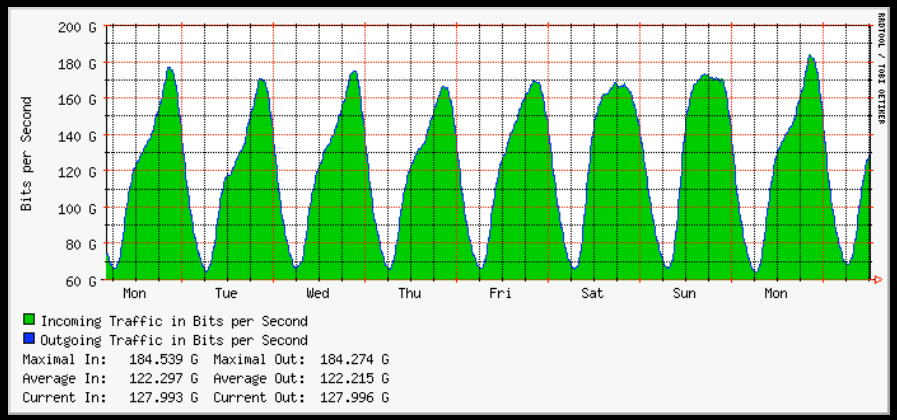

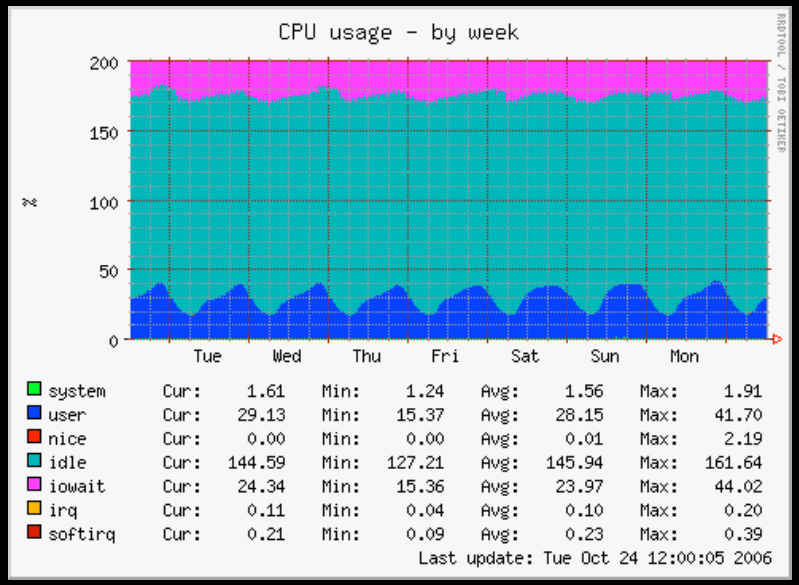

Performance Results

- I/O performance while writing data
  - Currently:
    - Writing ca. 40 000 RRD files in 8 seconds
  - High load tests:
    - Writing 130 000 RRD files in 27 seconds
- Max. at AMS-IX 160 000 conversations

Performance Results

- Foundry hardware
  - BigIron 15k
    - ASIC (Application-specific integrated circuit)
    - Switch CPU not affected
  - MG8 & RX*
    - Blade CPU affected

Performance Results

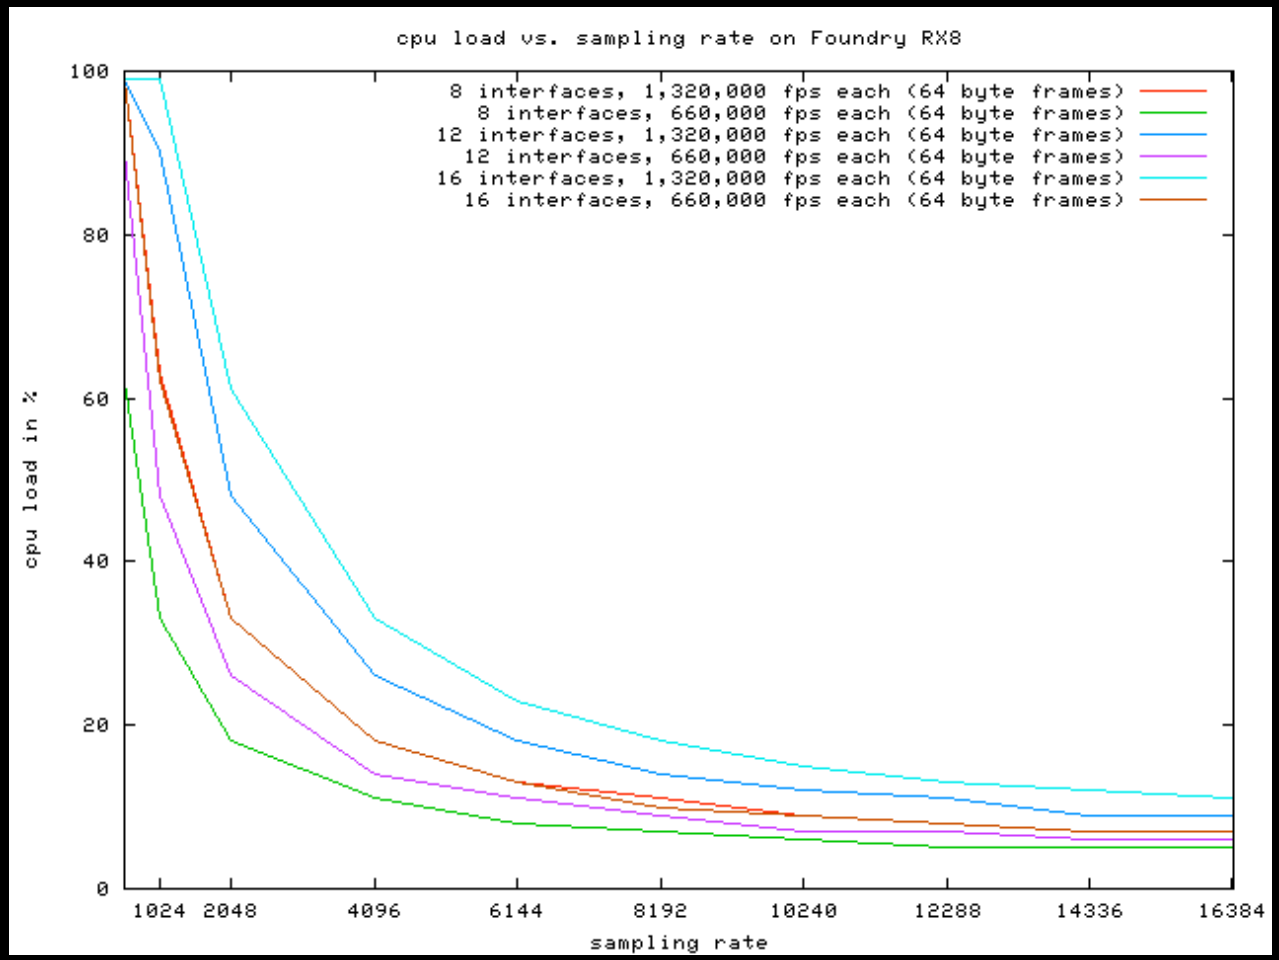

Software used at AMS-IX

- Analysis
  - Ether type graph – percentage of IPv4, IPv6, ARP and other
  - Total IPv6 traffic graph – in bps and pps
  - Member-2-Member analysis – in bps and pps

Privacy

- Statistical analysis
- Samples not saved after decoding
- Decoding only up to L2 (ethernet)
- More data not decoded by the software

Results

- Ether type - daily

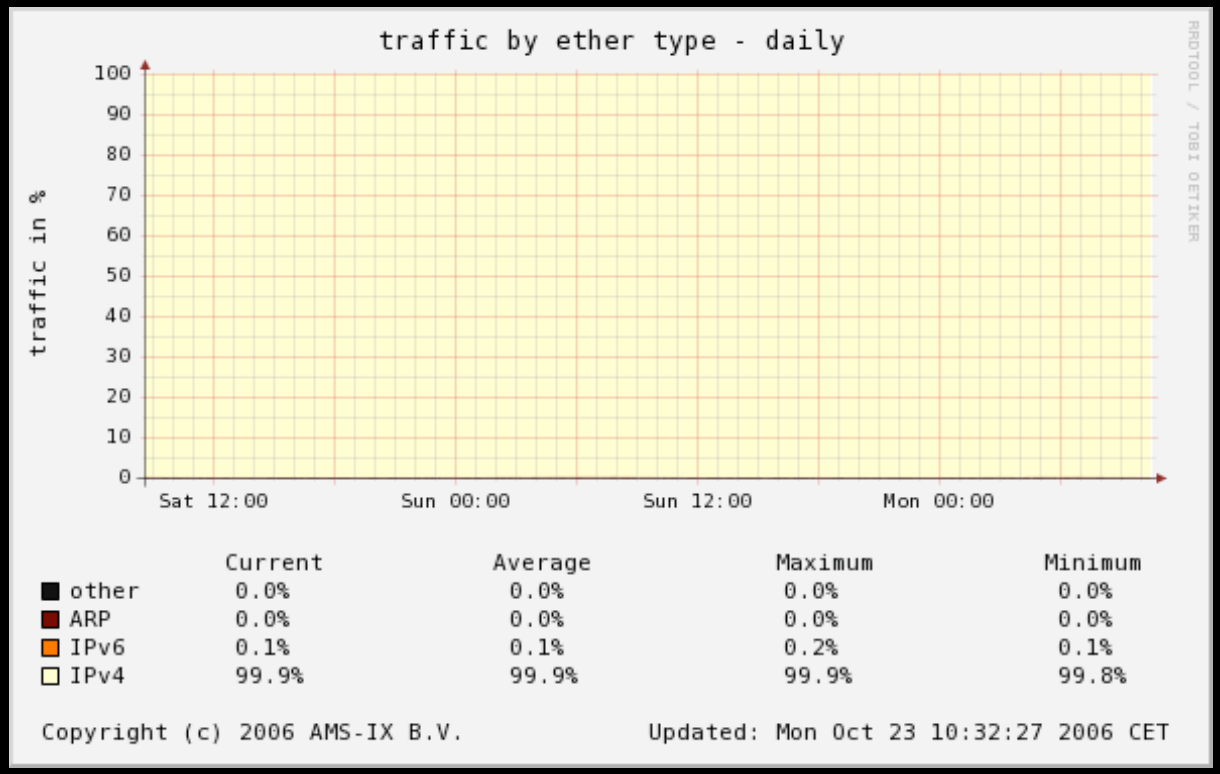

Results

- Ether type - monthly

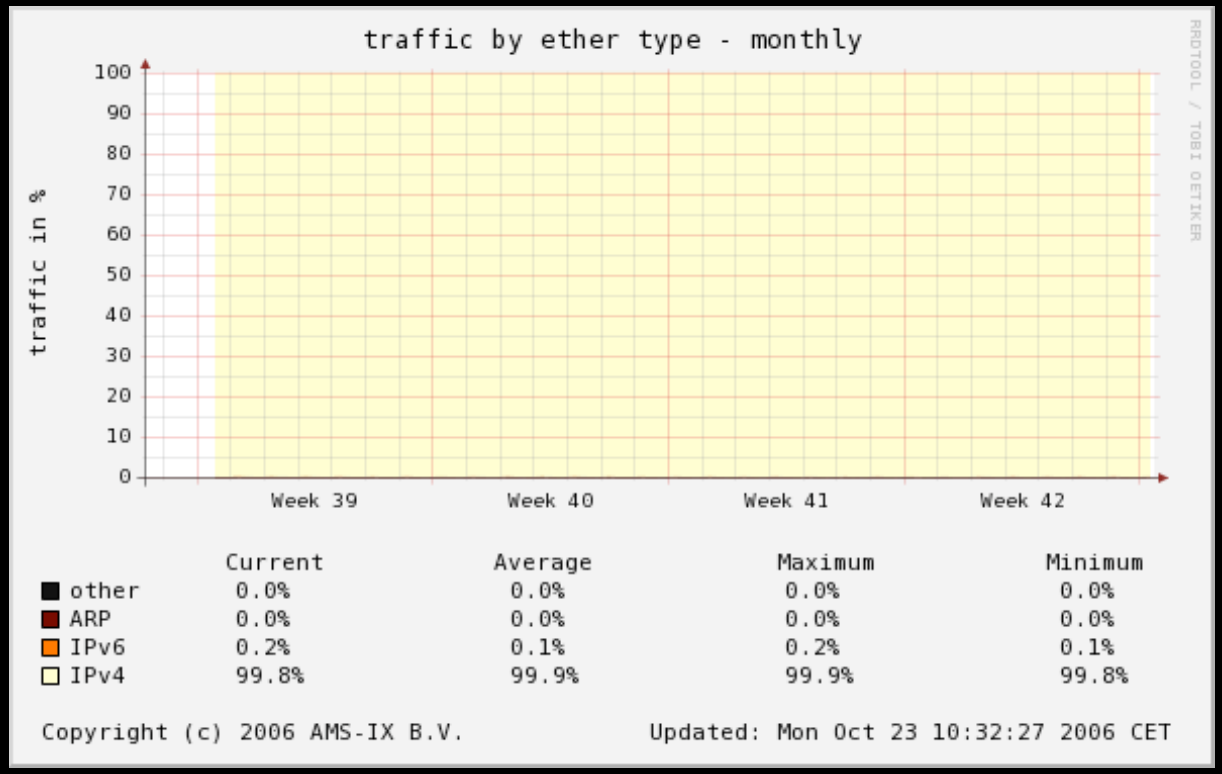

Results

- Total IPv6 traffic daily - bps

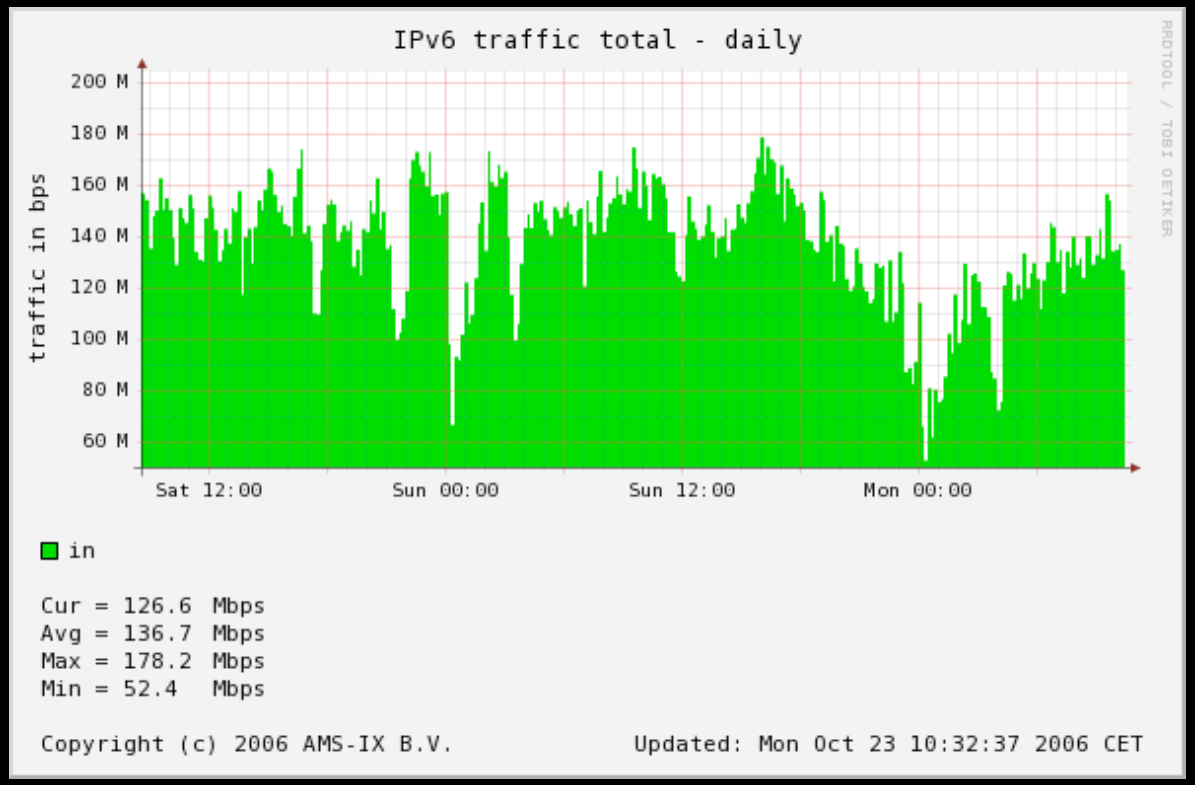

Results

- Total IPv6 traffic weekly - bps

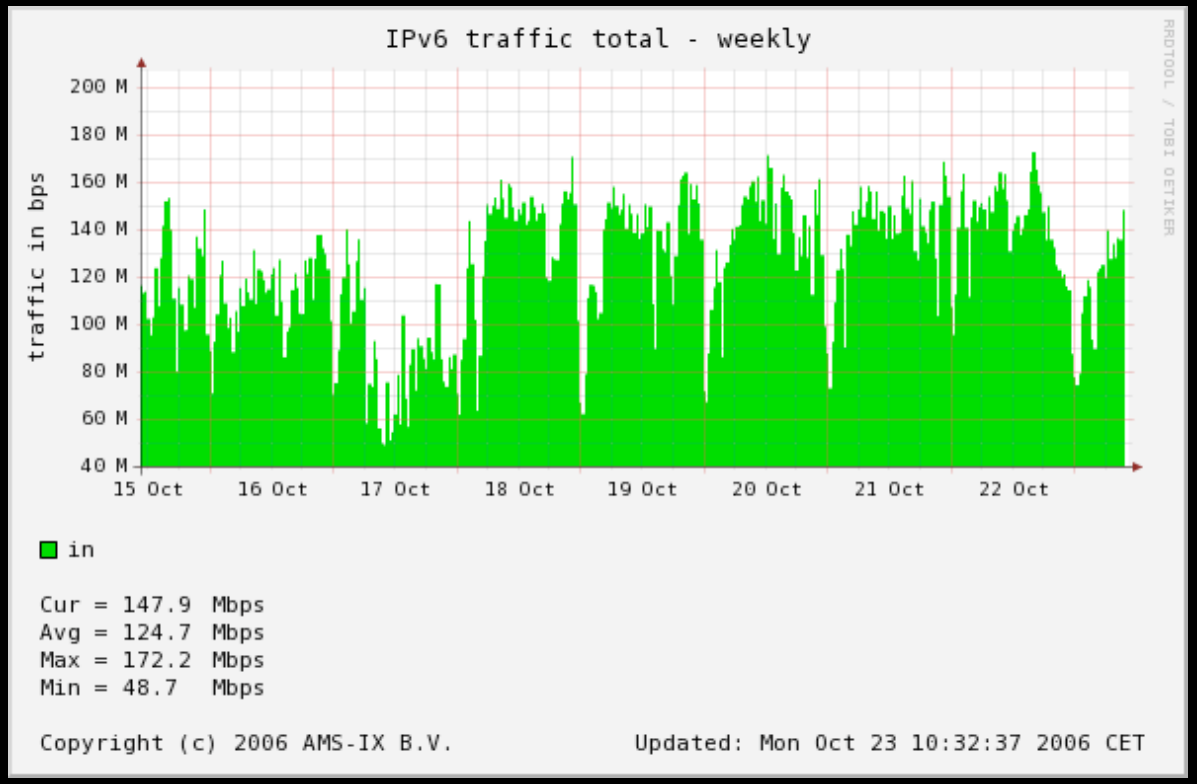

Results

- Total IPv6 traffic monthly - bps

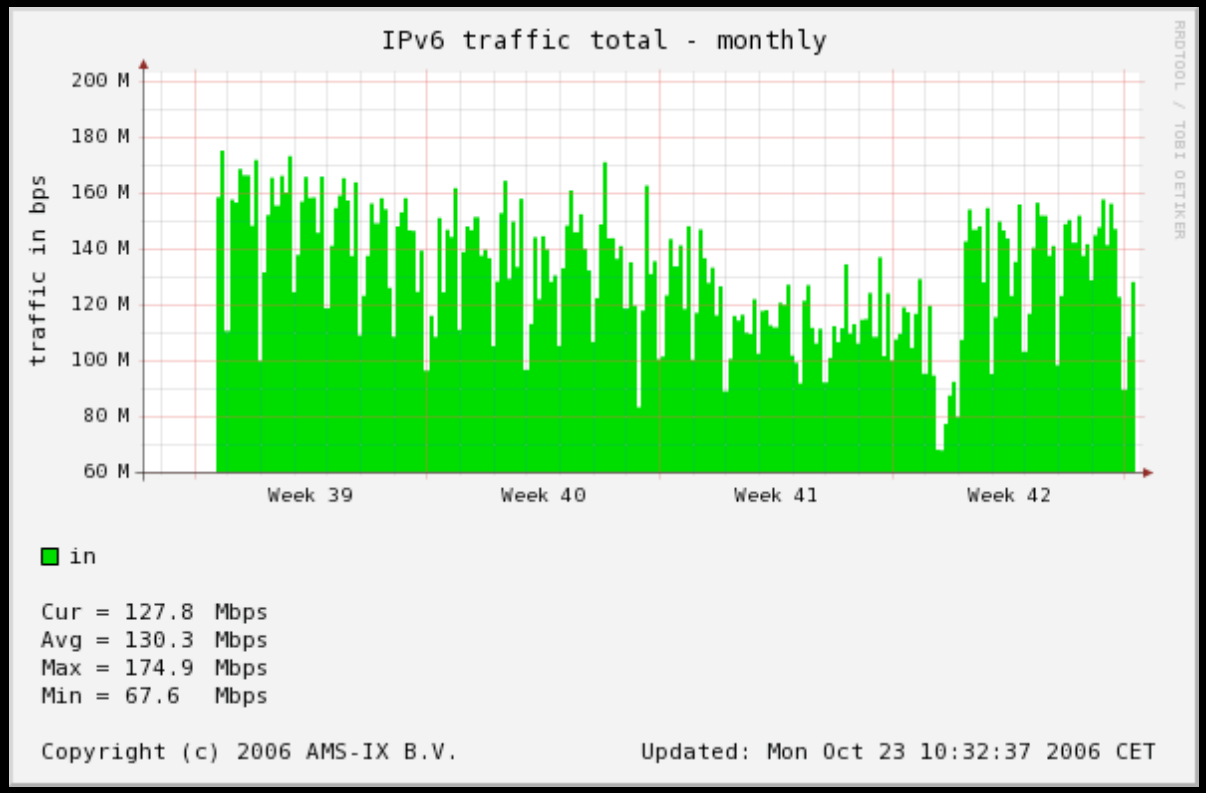

Results

- Member-2-Member traffic

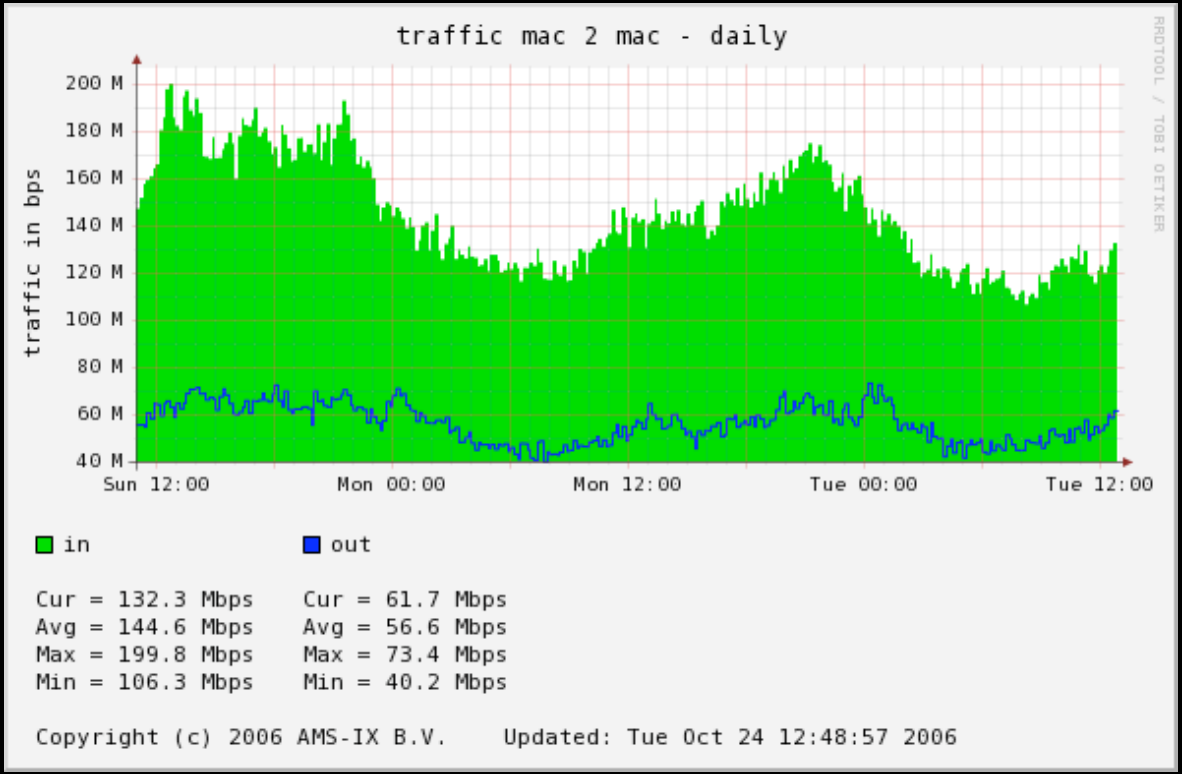

Results

- Member-2-Member traffic

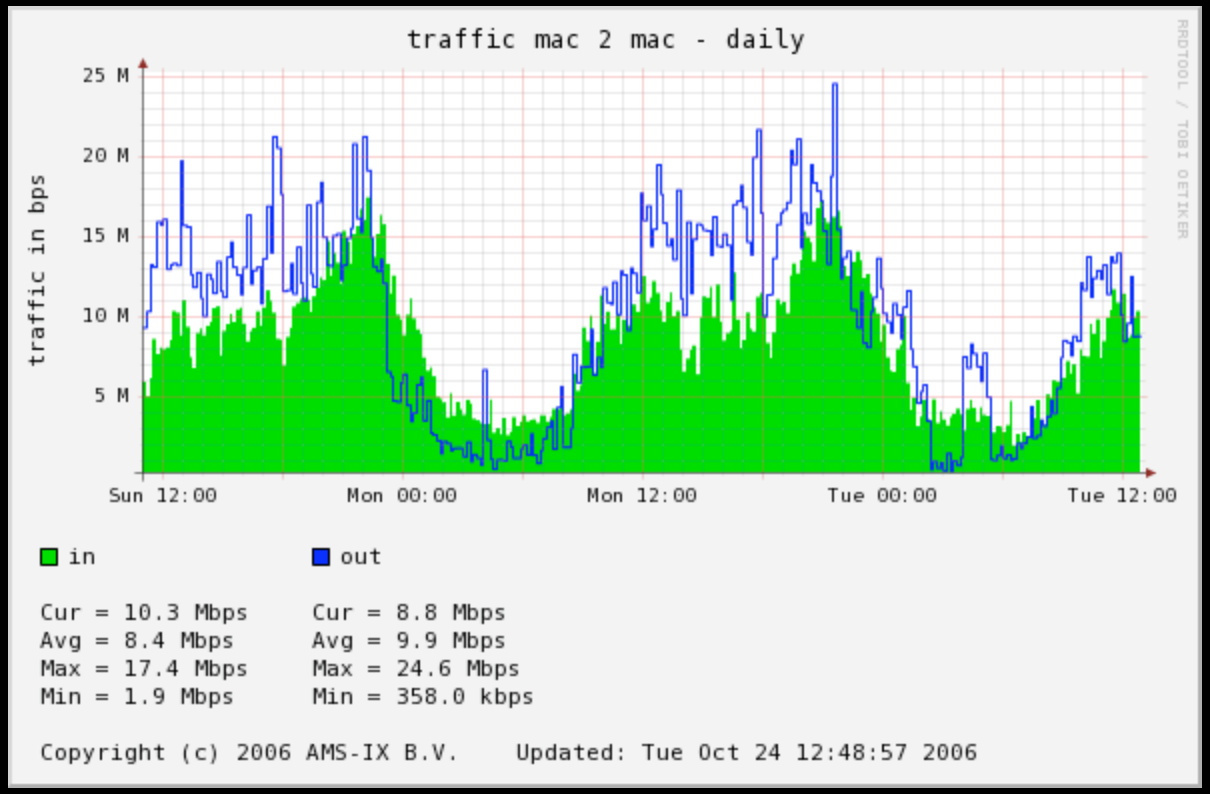

Results

- Member-2-Member traffic

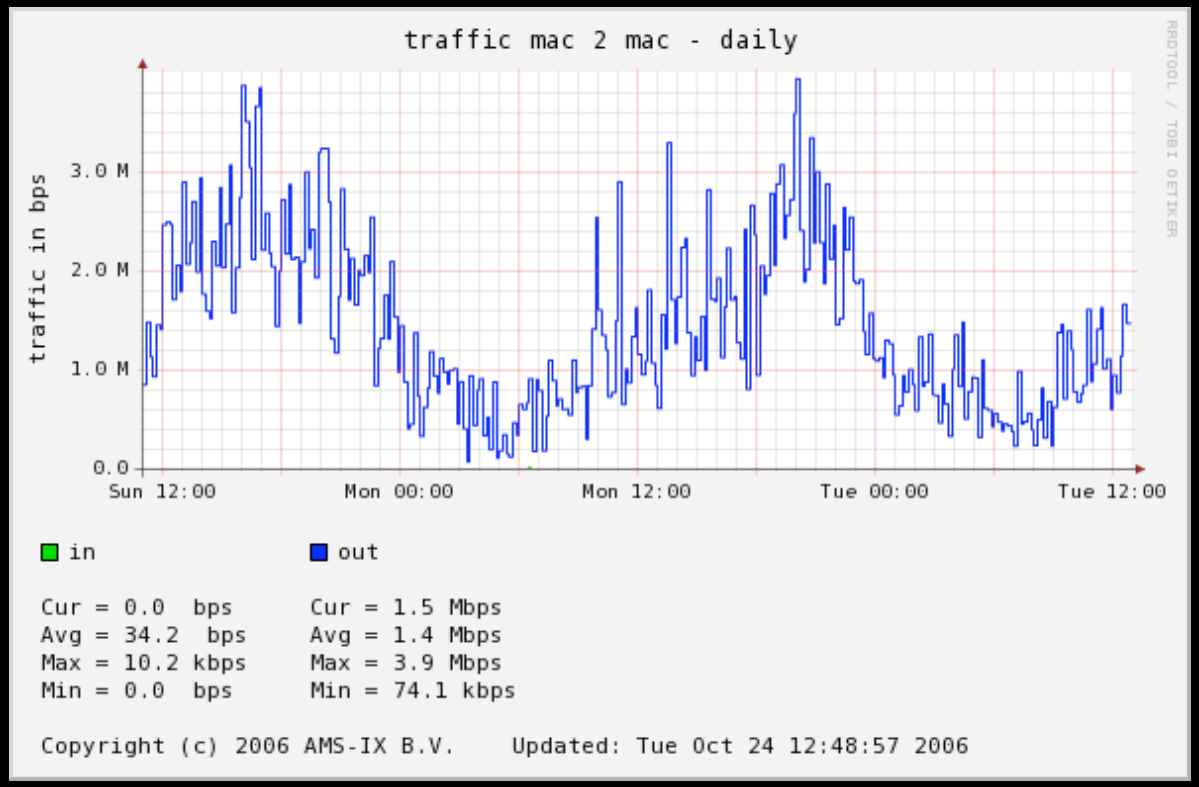

Future plans

- Use counter samples:
  - Separate interfaces
  - Aggregated links
  - Backbone links
  - Core network
  - ...

Questions ?

amsix
amsterdam internet exchange

Comments / Requests / Ideas:

- elisa.jasinska@ams-ix.net
- sflow@ams-ix.net