



# Adding QoS Attributes to the IPFIX Protocol

NORDUnet conference 2006  
Sept 28

Arne Øslebø  
[arne.oslebo@uninett.no](mailto:arne.oslebo@uninett.no)

# Outline

- Brief NetFlow/IPFIX introduction
- Motivation
- Implementation
  - Exporter
  - Collector
  - Presentation
- Performance
- Future work

# NetFlow

- Originally Cisco technology - 1996
- IETF is standardizing NetFlow in the IPFIX working group
- IPFIX definition of an IP network flow:
  - A set of IP packets passing an observation point in a network during a certain time interval. All packets belonging to a particular flow have a set of common properties.
- Flow Key
  - Each of the properties that are used for defining a flow
- Flow record
  - Measured properties of a flow


# Flow records

- NetFlow v5
  - Src&dst IP address
  - Next hop router's IP address
  - In&out interface index
  - Pkts and bytes in the flow
  - sysUptime at start and end of flow
  - TCP/UDP src and dst port number
  - Type of service
  - TCP flags
  - IP protocol
  - Src&dst AS number
  - Src&dst address prefix mask bits
- NetFlow v9
  - Configurable
- IPFIX
  - Based on NetFlow v9
  - Possible to add enterprise specific attributes

# Motivation

- Standard NetFlow can tell you things like:
  - which AS number you send the most traffic to
  - which IP address sends/receives the most traffic
  - +++++
- It can NOT tell you
  - who initiated the traffic
  - the quality of the traffic
    - retransmissions
    - jitter
    - burstiness
  - the type of application that generates traffic
- We want to extend IPFIX with enterprise specific attributes so that this information is available

# Framework



Stager user interface






<http://www.ist-lobster.org>



- Flow collector based on NERD
- Stager backend

- Passive monitoring card
- MAPI



# Exporter

- Supports: NetFlow v5, v9 and IPFIX
- Part of Monitoring API (MAPI)
  - <http://mapi.uninett.no>
- New attributes:
  - pktLenHistogram, pktDistHistogram
  - pktPayload
  - pkt[Dist/Length][Var/Sum/SumQ]
  - direction
  - reordered
  - rtpJitter, rtpLostFraction, rtpLostPackets, rtpSequenceCycles
  - maxRate[1sec/100ms/10ms/1ms]
  - minRate[1sec/100ms/10ms/1ms]
  - service

# Measurement probes



# Collector

- Modified version of NERD
  - <http://www.nerdd.org>
- Stores flow records to file
- FlowStat generates high level reports that are inserted into the Stager database.

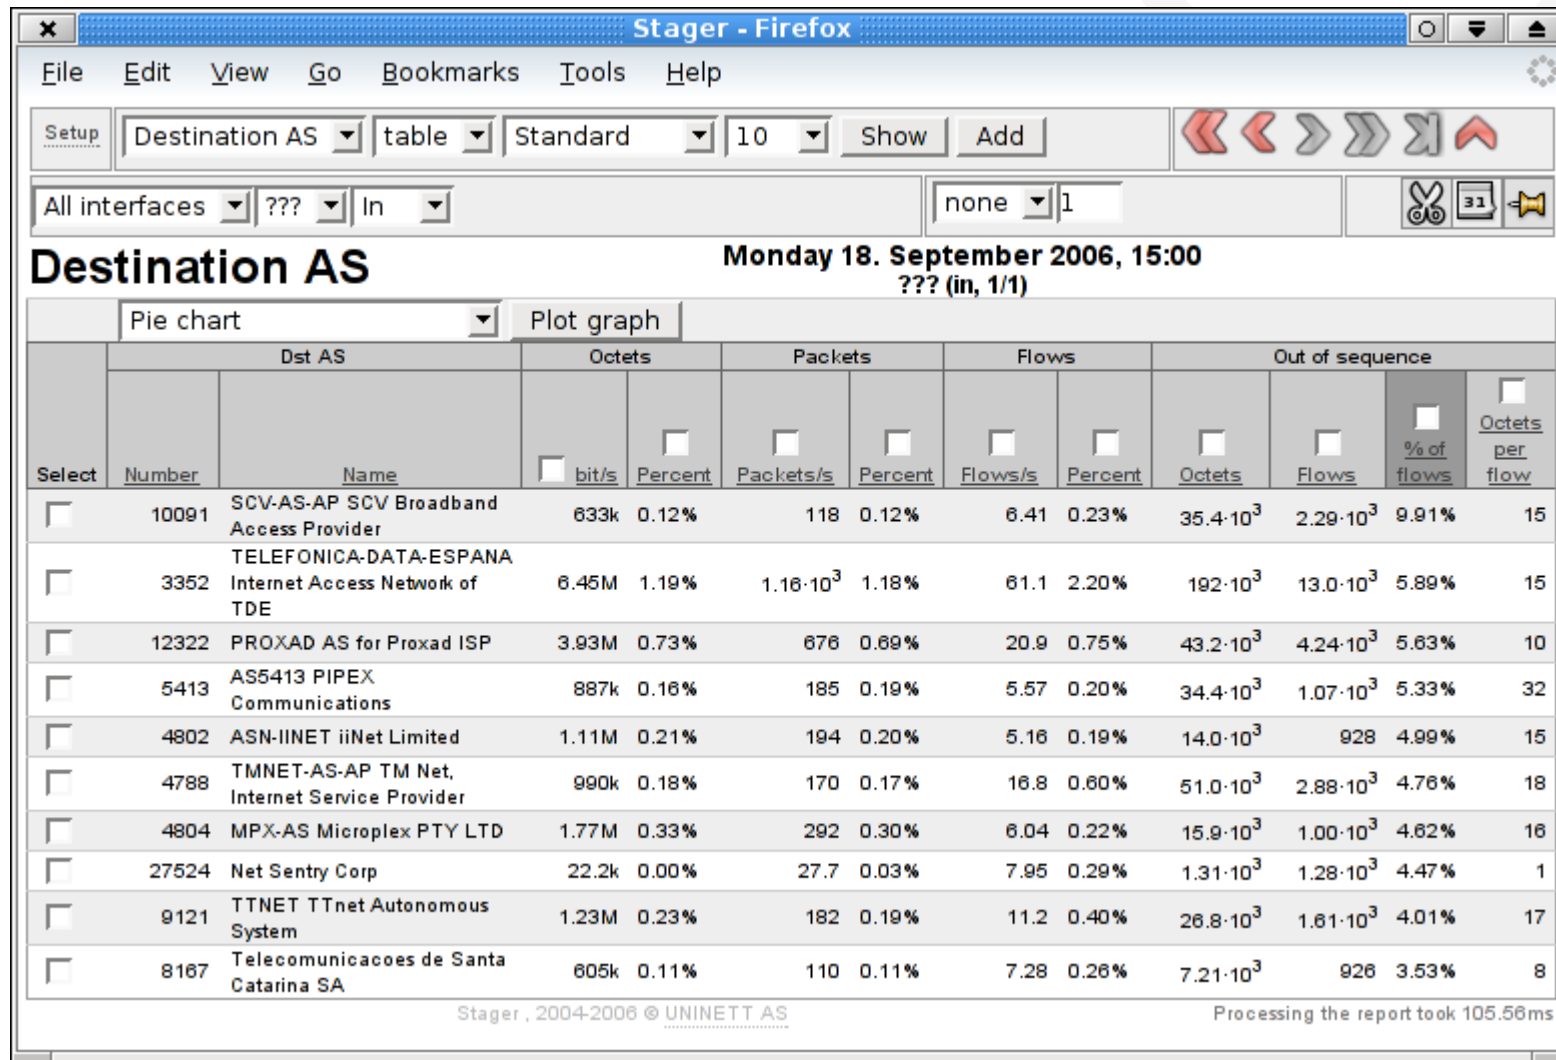
9

```
FlowStat -o "2 desc" -f "packets>500" -s "dst_as sum(octets) avg(hpktdst)"  
# dst_as sum(octets) avg(hpktdst)  
0, 800721047, 28_0c_08_0c_4c_27_08_04_0c_d3_,  
15659, 719072910, 1a_06_03_04_20_07_04_02_08_f2_,  
8642, 615117228, 25_0c_07_0c_47_16_07_0a_10_dd_,  
3301, 592188599, 12_04_05_0a_3f_0e_02_01_0b_ef_,  
2119, 446154823, 1d_06_03_06_2a_07_02_01_04_f0_,
```

# Stager – user interface

- A web-based tool for aggregating and presenting most types of network statistics
- GPL
- 140+ people on a public mailinglist
- Easy to use web frontend
  - Text based reports and graphs
- Easy to add new reports
  - Templates and plugins
- Statistical functions
- Access control
  - Observation points and reports
- Scalable
  - Handle large volumes of data and store years of statistics

# Destination AS report



11

# Destination AS report (2)

Stager - Firefox

File Edit View Go Bookmarks Tools Help

Setup Destination AS table Connections 10 Show Add

All interfaces ??? In none 1

## Destination AS

Monday 18. September 2006, 15:00  
??? (in, 1/1)

Pie chart Plot graph

Select	Dst AS		Traffic		Requests		Responses	
	Number	Name	Octets	Flows	Percent	Per second	Percent	Per second
<input type="checkbox"/>	0	0	15.4G	480·10 <sup>3</sup>	26.98%	36	8.80%	11.7
<input type="checkbox"/>	15659	NEXTGENTEL NEXTGENTEL Autonomous System	15.1G	119·10 <sup>3</sup>	13.88%	4.6	34.76%	11.5
<input type="checkbox"/>	8642	B2 B2 Bredband AB (publ)	11.2G	73.4·10 <sup>3</sup>	12.12%	2.47	9.51%	1.94
<input type="checkbox"/>	2119	TELENOR-NEXTEL Telenor Internet Access	8.43G	135·10 <sup>3</sup>	15.05%	5.65	26.88%	10.1
<input type="checkbox"/>	3301	TELIA NET-SWEDEN TeliaNet Sweden	8.05G	130·10 <sup>3</sup>	30.64%	11.1	7.80%	2.81
<input type="checkbox"/>	3320	DTAG Deutsche Telekom AG	5.91G	158·10 <sup>3</sup>	4.31%	1.89	5.61%	2.46
<input type="checkbox"/>	6830	UPC UPC Broadband	4.64G	92.7·10 <sup>3</sup>	6.77%	1.74	7.31%	1.88
<input type="checkbox"/>	7132	SBC Internet Services	4.24G	110·10 <sup>3</sup>	6.82%	2.08	6.05%	1.85
<input type="checkbox"/>	8394	ALFANETT Alfannett Autonomous System	3.61G	28.8·10 <sup>3</sup>	7.67%	0.613	33.54%	2.68
<input type="checkbox"/>	3215	AS3215 France Telecom Transpac	3.55G	93.6·10 <sup>3</sup>	4.80%	1.25	10.18%	2.64

Stager, 2004-2006 © UNINETT AS Processing the report took 79.79ms

12

# Destination IP report

Stager - Firefox

File Edit View Go Bookmarks Tools Help

Setup Destination IP table Standard 20 Show Add

All interfaces ??? In none 1

## Destination IP

Monday 18. September 2006, 15:00  
??? (in, 1/1)

Pie chart Plot graph

Select	Dst IP	Octets	Rate 1 second		Rate 100 milliseconds		Rate 10 milliseconds	
		Total	Max	Min	Max	Min	Max	Min
<input type="checkbox"/>	w.x.y.z	18.8M	10.5M	780k	15.8M	0	53.8M	0
<input type="checkbox"/>	w.x.y.z	8.91M	5.21M	583k	5.37M	0	32.4M	0
<input type="checkbox"/>	w.x.y.z	7.75M	195k	75.3k	1.33M	163k	4.72M	5.78k
<input type="checkbox"/>	w.x.y.z	8.79M	143k	140k	937k	731k	7.62M	315k
<input type="checkbox"/>	w.x.y.z	7.48M	45.9k	2.38k	143k	1.88k	1.22M	7.34k

Stager, 2004-2006 © UNINETT AS Processing the report took 131.8ms

13

# Performance - exporter

- Depends on included attributes in the flow records
- Most common attributes
  - Tested on OC48 without problems

Element	Instructions
MAPI	200
Normal attributes	680
HIST_PKT_LEN	30
HIST_PKT_DIST	495
VAR of DIST/LEN	100
MAX/MIN RATE	300
SERVICE	1760

# Performance - collector

- Collecting and storing to disk is no problem
- Current bottleneck:
  - Processing data to generate high level reports
  - >1 hour to process 1 hour of data
- Main problem
  - Not possible to generate multiple reports in one pass
  - Some reports demands multiple passes

# Performance - Stager

- Highly scalable
- Our NetFlow setup:
  - 27 routers
  - 207 interfaces
  - >30Gb of raw Netflow data every day
  - 400.000 new entries in the db every hour
  - >450 millions entries in a single table
  - >700Gb database size

16

	<b>PC1</b>	<b>PC2</b>	<b>PC3</b>	<b>Total</b>
<b>Netflow size</b>	537MB	161MB	399MB	1097MB
<b>Sequentially</b>	9min 17s	3min 8s	5min 51s	18min 16s
<b>No insert in DB</b>	7min 11s	1min 48s	4min 52s	13min 51s
<b>Simultaneously</b>	9min 21s	3min 7s	5min 56s	18min 24s
<b># of new DB entries</b>	164702	69549	184706	418957
<b># of entries/second</b>	295.69	369.94	526.23	382.26

# Future work

- Verify correctness of attributes
- Improve performance of FlowStat
  - Generate multiple reports on one pass over raw NetFlow files
- Full deployment on all measurement probes
  - Gain more experience to see which attributes are the most useful

# Software availability

- <http://software.uninett.no>
- E-mail contact:
  - [arne.oslebo@uninett.no](mailto:arne.oslebo@uninett.no)

18



<http://www.ist-lobster.org>