Traffic Accounting

Elisa Jasinska
<ejas@microsoft.com>

Welcome

- Aaron Hughes (6connect)
- Paolo Lucente (Cariden/pmacct.net)
- Brent Van Dussen (Limelight Networks)
- Arien Vijn (AMS-IX)

Network Traffic Accounting

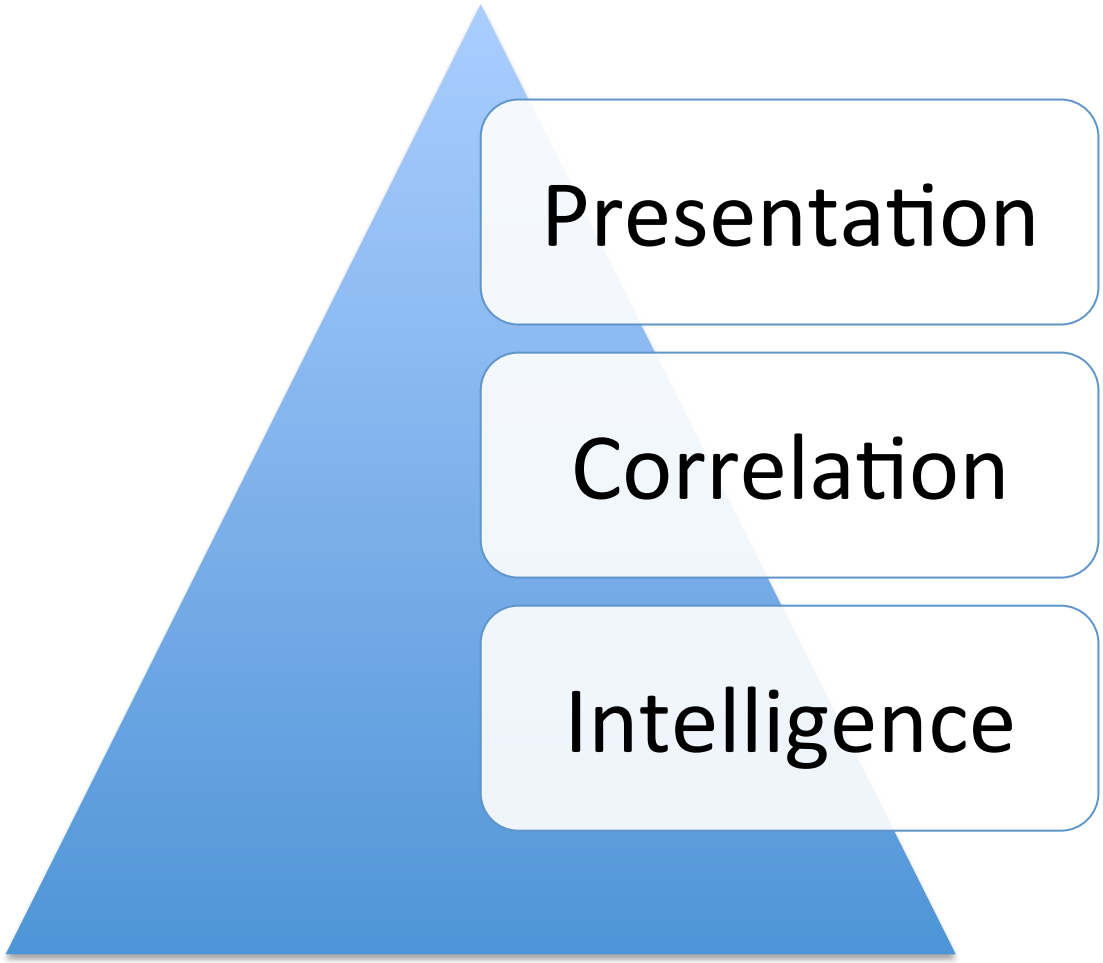

Tools

- SNMP
  - Cacti
  - MRTG
  - Cricket
- Flow
  - Appliances
  - pmacct
  - Minimalistic parsing tools
    (sflowtool, Net::sFlow, nfdump)

Problems

- Lots of effort into collection and visualization
- Development resources needed: heavy sysadmin, dbadmin, integration work
- Scalability of DB/storage

Introduction

Aaron Hughes
6Connect

Traditional SNMP

- SNMP
  - 95th percentiles
  - Interfaces getting larger and 32bit -> 64bit counters / rollover problems
  - Not a very good solution for billing
  - RRD style storage and no real way to go back for old time periods without ave kicking in
  - Differentiating internal and external traffic is, in many cases, hard to do (e.g. replication vs. transit)

Traditional DPI

- Not many things which can tap correctly in aggregation
- Limited to mirror ports, endace cards etc.
- Copying traffic to central analytics location is very expensive
- Triple data volume for 1:1 analytics
- I/O limitations of storing this volume of data
- Analytics tools take a LONG time for large volume
- Privacy related issues

Traditional Flow

- Multiple versions
- Vendors treat flow data differently
- Managing e-gress i-gress interface tagging is challenging
- Desired sampling rates cost processor
- Flow processor usage can amply DDoS attacks

Traditional Visualization

- Rrdtool
- Cacti
- MRTG
- Home Grown

Traditional and Emerging Trends for SPs

Paolo Lucente

Cariden Technologies / pmacct.net

Traffic accounting: traditional trends

- Security, triggering alarms
- Customer billing
- Historical trending
- Monitor customer quotas, fair-usage policies
- Monitor own QoS functionality
- Capacity planning, Traffic engineering
- Analyze internet peering & transit:
  - Ie. Monitor traffic ratios, monitor how much traffic is exchanged with whom, study how to reduce IP transit bill, detect revenue leaks, etc.

Traffic accounting: emerging trends

- In-depth analysis of IP transit (BI):
  - Customer profitability
  - OTT caches behaviour
  - Peering traffic distance (ie. bit miles calculations)
- SDN:
  - Holistic view in a central controller can help increasing quality of routing decisions
- Mobile:
  - Increasing interest from mobile operators as their infrastructures get more and more into IP

Traffic Accounting at Limelight

Brent Van Dussen
Limelight Networks

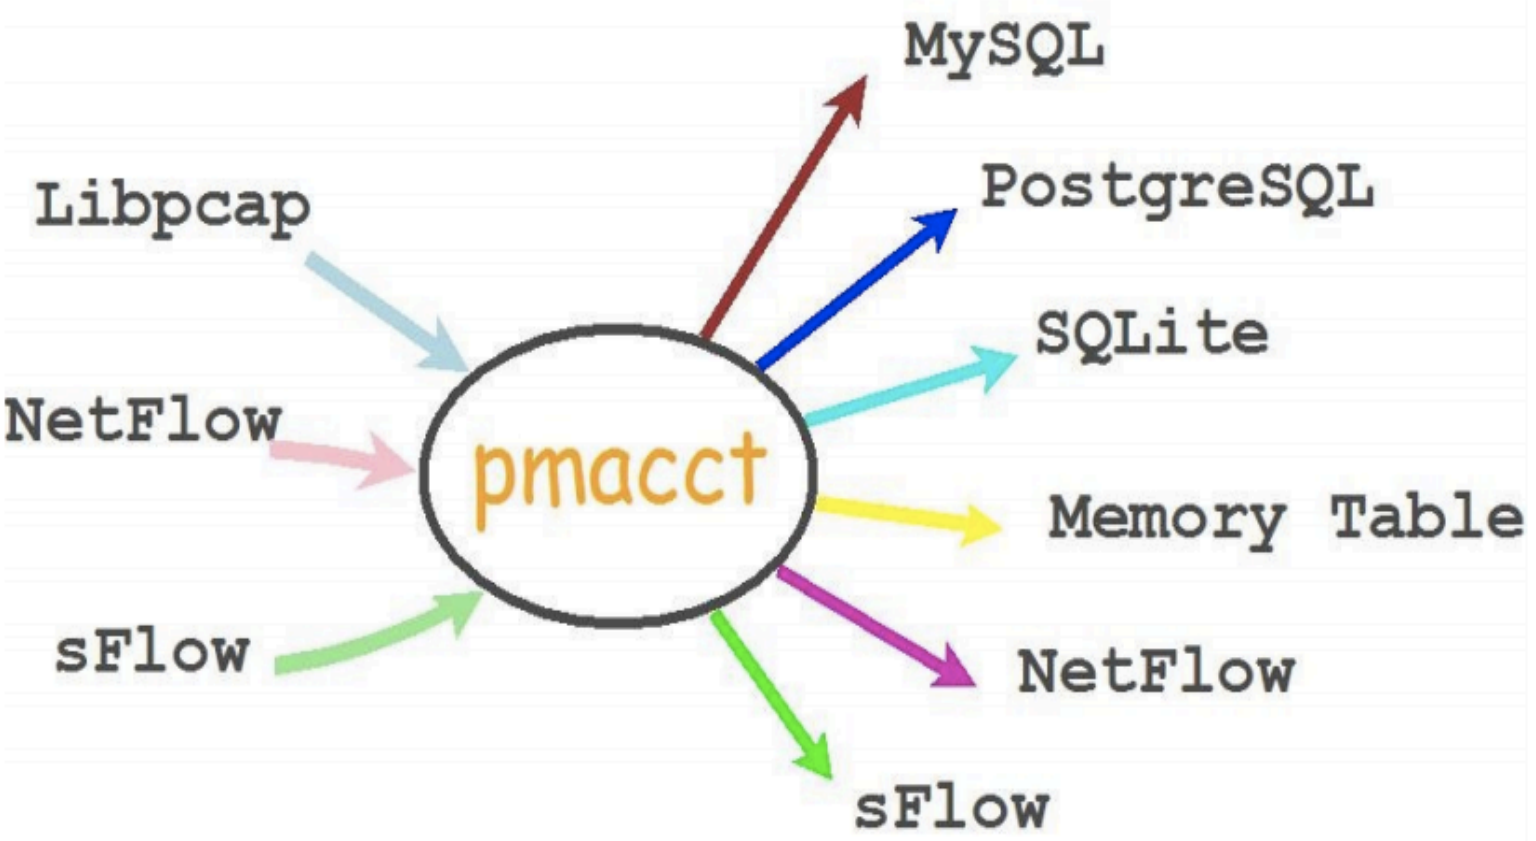
MySQL
PostgreSQL
Libpcap
SQLite
NetFlow
pmacct
Memory Table
sFlow
NetFlow
sFlow

pmacct (cont.)

- Open Source flow collection/aggregation tool
- Supports sflow/netflow/ipfix
- Flexible backend storage options
- www.pmacct.net

Collector Architecture

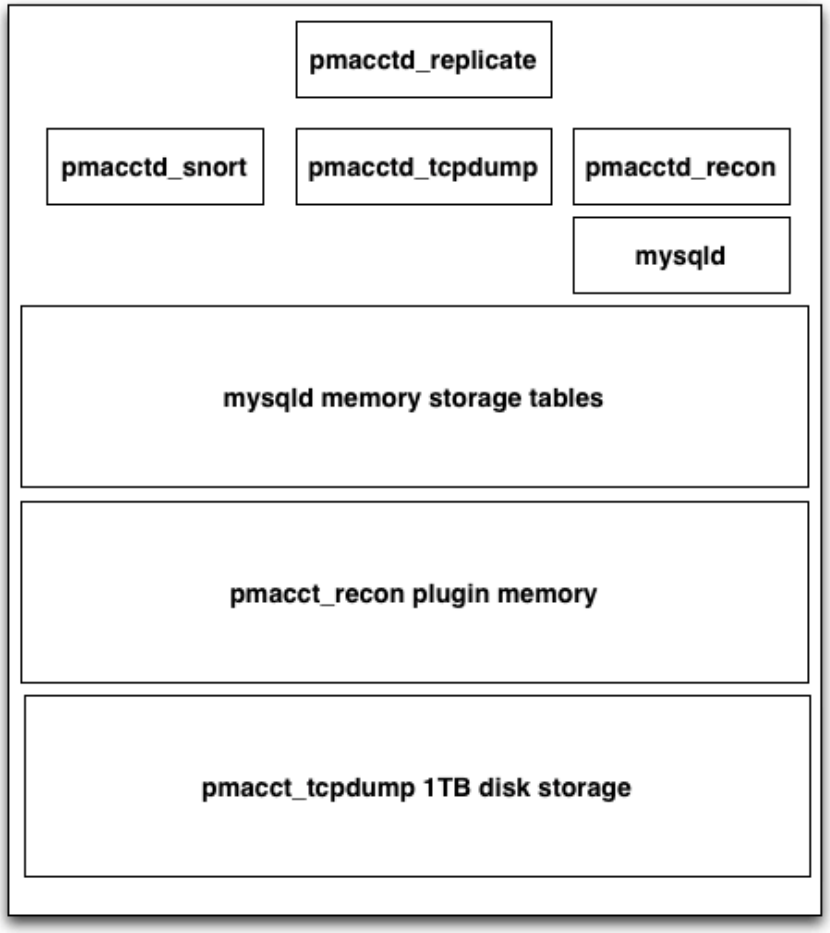

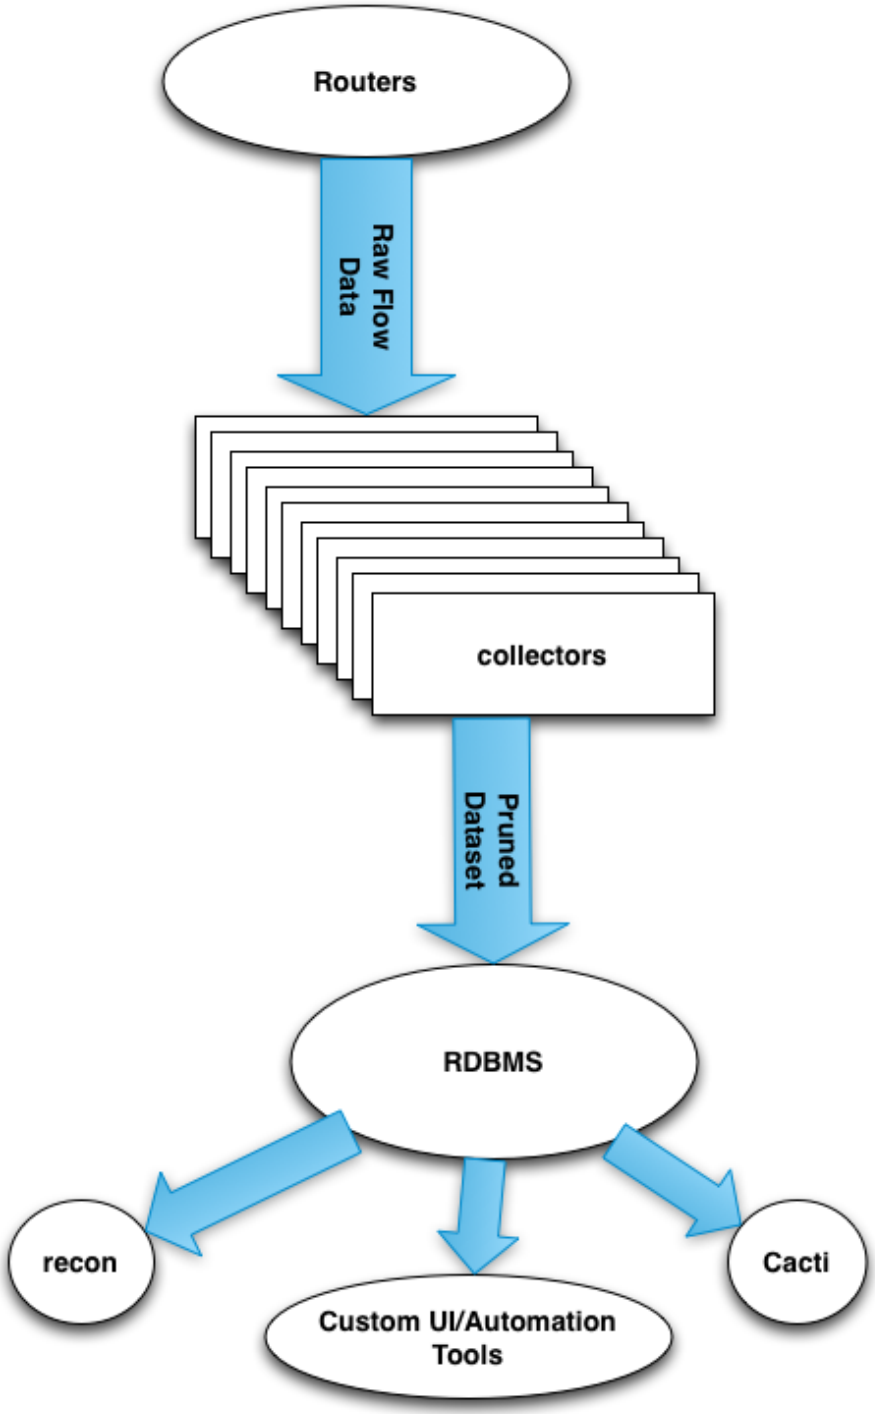
Routers
Raw Flow Data
collectors
Pruned Dataset
RDBMS
recon
Custom UI/Automation Tools
Cacti

Neat Data!

- Backend RDBMS makes it easy to query for specific data
- Ipv4/ipv6 breakdown per interface
- Insight into backbone/egress flows broken down by offnet/onnet
- Visualize traffic going places that it should not go
- Command line utilities to extract near real time as-path/prefix information per interface broken down by pps/bps
- Snort IDS logs to splunk for additional security data/alerting
- Tcpdump saves raw flow data to disk for post-mortem signature analysis/forensics

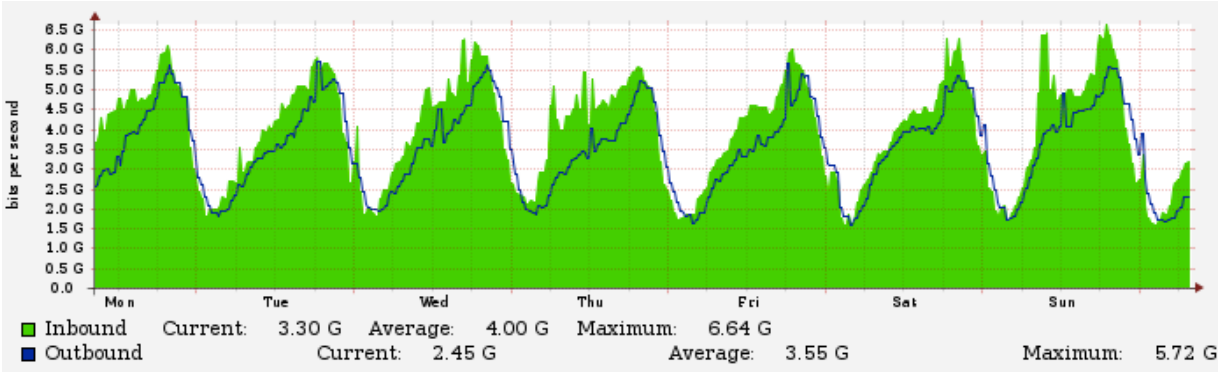
bits per second
6.5 G
6.0 G
5.5 G
5.0 G
4.5 G
4.0 G
3.5 G
3.0 G
2.5 G
2.0 G
1.5 G
1.0 G
0.5 G
0.0
Mon
Tue
Wed
Thu
Fri
Sat
Sun
Inbound Current: 3.30 G Average: 4.00 G Maximum: 6.64 G
Outbound Current: 2.45 G Average: 3.55 G Maximum: 5.72 G

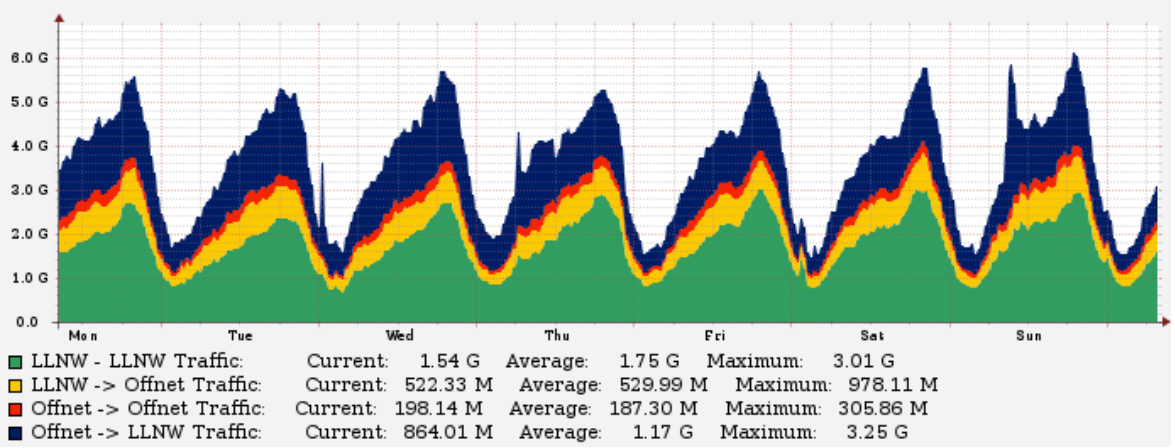
6.0 G
5.0 G
4.0 G
3.0 G
2.0 G
1.0 G
0.0
Mon
Tue
Wed
Thu
Fri
Sat
Sun
LLNW - LLNW Traffic: Current: 1.54 G Average: 1.75 G Maximum: 3.01 G
LLNW -> Offnet Traffic: Current: 522.33 M Average: 529.99 M Maximum: 978.11 M
Offnet -> Offnet Traffic: Current: 198.14 M Average: 187.30 M Maximum: 305.86 M
Offnet -> LLNW Traffic: Current: 864.01 M Average: 1.17 G Maximum: 3.25 G

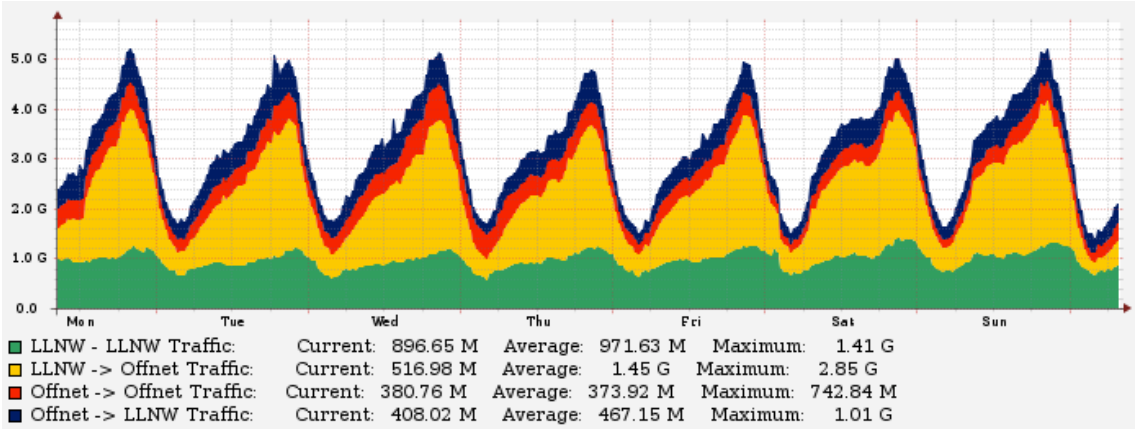
5.0 G
4.0 G
3.0 G
2.0 G
1.0 G
0.0
Mon
Tue
Wed
Thu
Fri
Sat
Sun
LLNW - LLNW Traffic: Current: 896.65 M Average: 971.63 M Maximum: 1.41 G
LLNW -> Offnet Traffic: Current: 516.98 M Average: 1.45 G Maximum: 2.85 G
Offnet -> Offnet Traffic: Current: 380.76 M Average: 373.92 M Maximum: 742.84 M
Offnet -> LLNW Traffic: Current: 408.02 M Average: 467.15 M Maximum: 1.01 G

Go IPv6!

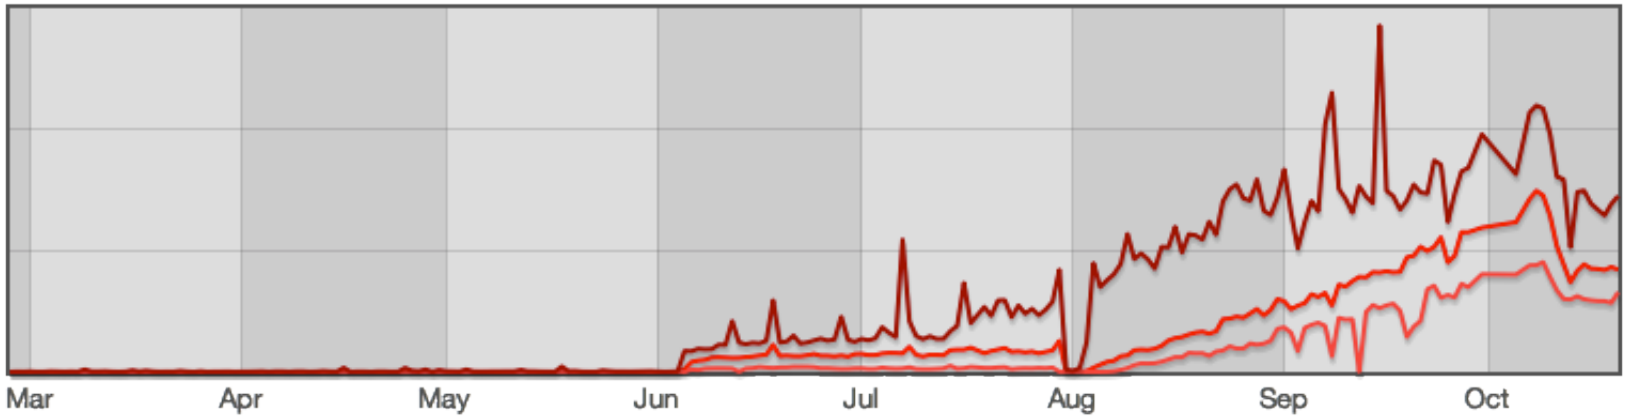

recon

bvd@cholla:~$ recon –l <router hostname> -p 15/1 -ad
****
##
TIME: 2012-10-23 10:30:00 to 2012-10-23 10:34:59
##

AS DST TRAFFIC (Mbit/s) AVG LOCATION IFACE
##
xxx1 820.0 <router hostname> 561
xxx7 328.0 <router hostname> 561
xxxx8 255.0 <router hostname> 561
xx1 241.0 <router hostname> 561
xxx3 230.0 <router hostname> 561
xxxx2 211.0 <router hostname> 561

Future

- Nosql development underway
- MongoDB? Cassandra?
- Parallelization of insert/select for improved performance and redundancy
- Automated traffic management platform

Traffic Accounting at AMS-IX

Arien Vijn
Amsterdam Internet Exchange

Traffic Accounting

- Pricing is NOT based on traffic accounting.
  - Fixed price per port per month.
- Traffic graphs are there for the following reasons:
  - Service to our members/customers.
  - Traffic growth analyses.
  - Troubleshooting.

What do we have?

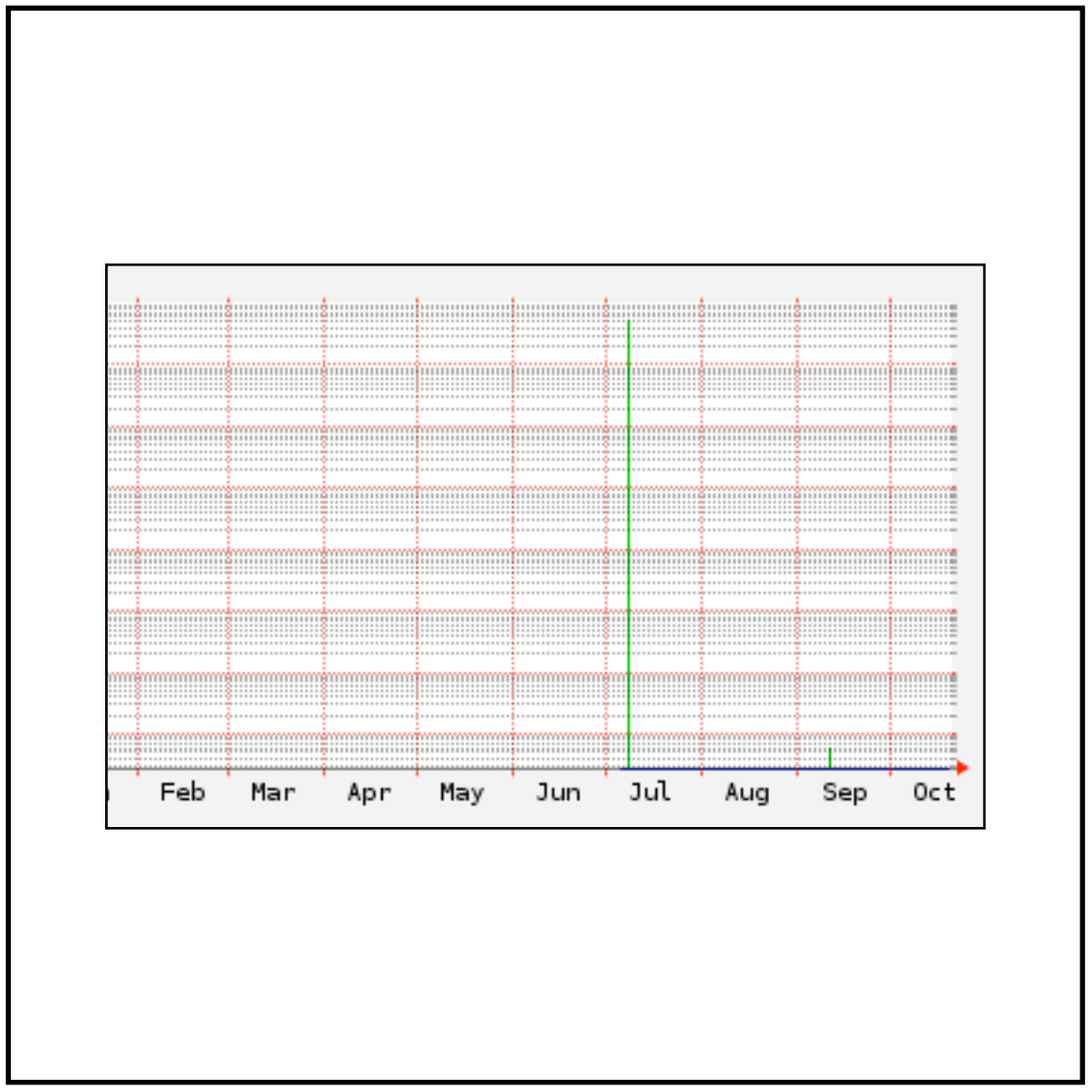

- Customer/Member ports:
  - SNMP counters:
    - Bits, Frames, and Errors.
  - Presented as time series data (RRDtool).
    - Web portal
    - Optional logarithmic y-axis.
  - Public traffic graph is the sum of all member ports.

What do we have?

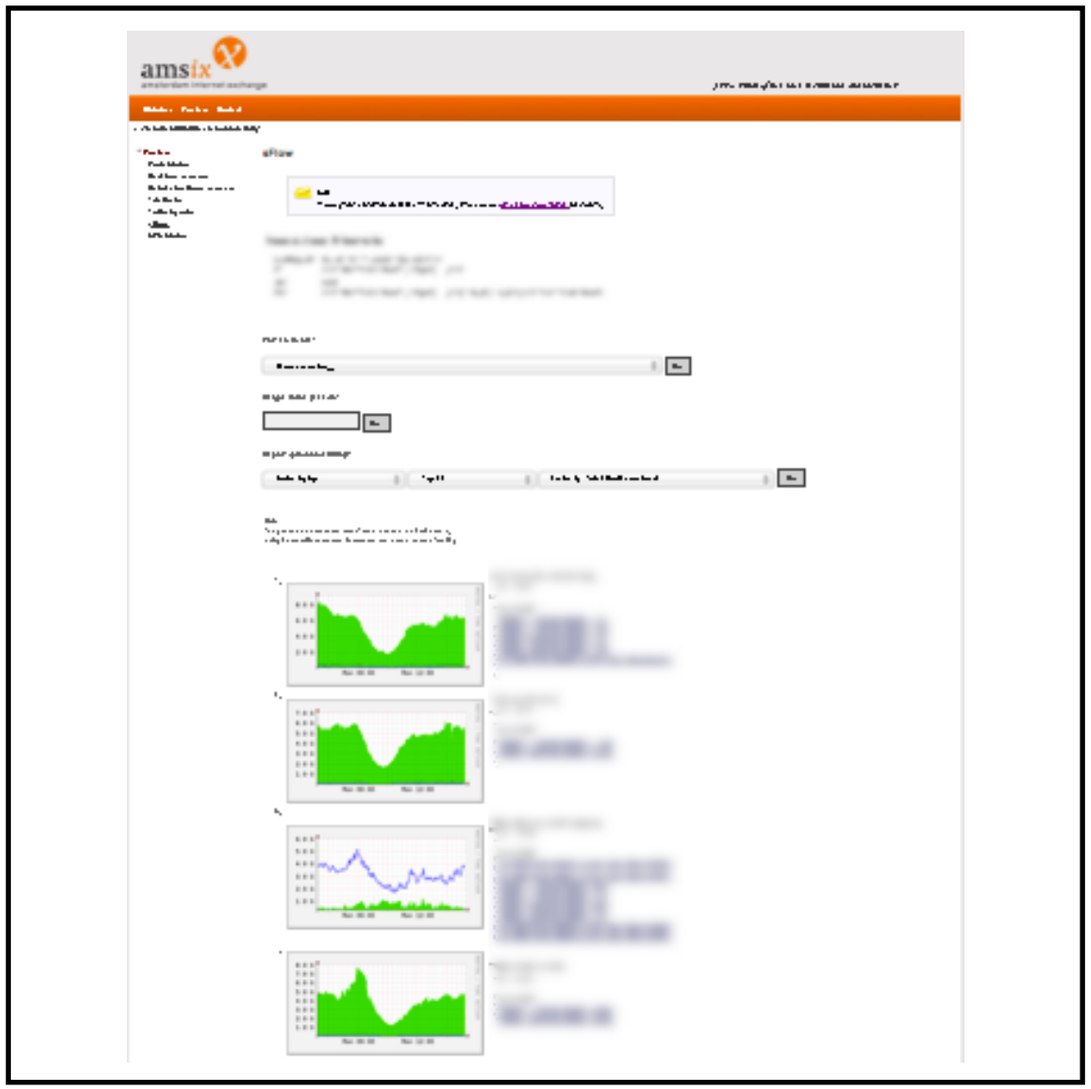

- sFlow:
  - Net::sFlow (PERL).
  - Only source and destination MAC addresses and Ethertype.
- Presented as time series data (RRDtool).
  - Web portal allows to see traffic per peer.
  - Very useful tool for traffic engineering and peering decisions.

What do we have?

- Key Performance Identifiers
  - Delay.
  - Jitter.
  - Frameloss.
- Active monitoring
- Presented as time series data that is updated every 5 second (SynLeaf).

What do we have?

From: noc@ams-ix.net (AMS-IX NOC)
Subject: RRD Daily Summary Report
Date: October 22, 2012(43) 8:12:20 PM CDT
To: noc@ams-ix.net

Ports with more than 2% error samples between 201210220311 and 201210230311:

- core-xxx-yyy-bb:stub-zzz-000 (AMS-IX): https://stats.ams-ix.net/

Ports with more than 95% average bandwidth use between 201210220311 and 201210230311:

- link-12345678-105-002 (ErrCom): 97.11%/2.72%
https://stats.ams-ix.net/

- Warning emails
  - Errors.
  - Over-utilized links and ports.

What do we have?

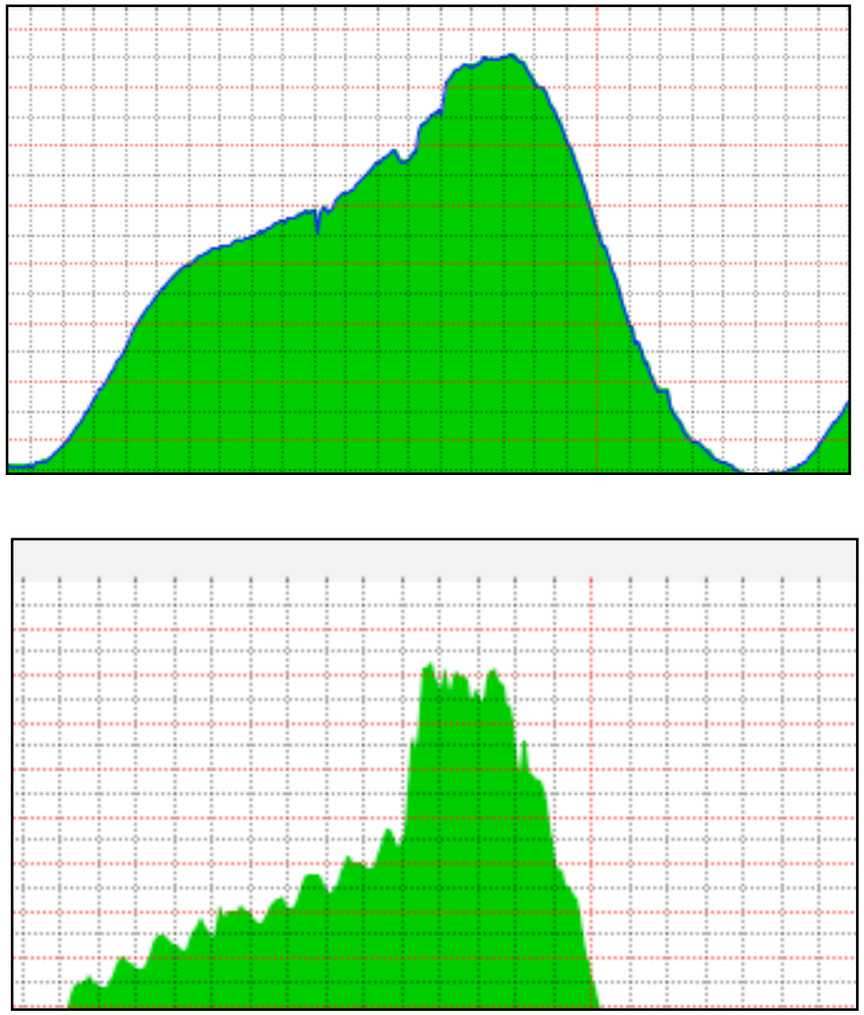

- Correlation tool
  - Compares a selected period with historical data.
  - Find anomalies.

What are the issues?

- Huge number of RRD files
  - File system database.
  - Digging through all files to do any analyses.
  - Updating huge numbers of files every 5 minutes.
  - Slow.

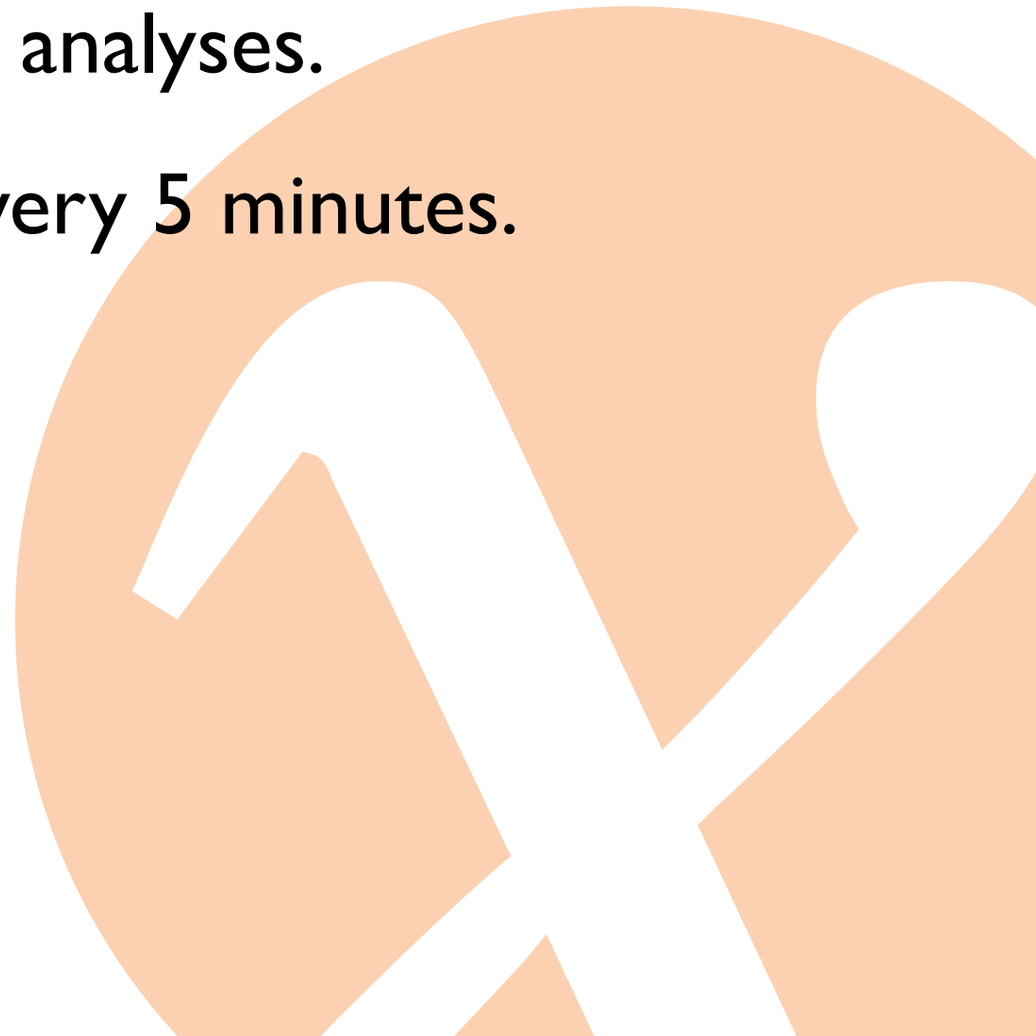

What would we like?

(i think)

- sFlow counters
- Multi dimensional RRD-like data structure
  - Fast queries
  - Fast updates
  - Scalable
    - Multi host system?
    - Erlang anybody?

Conclusion

- Endless possibilities unlocked by current innovation
  - Correlation of flow data with IGP/EGP
- There are alternative options
- The surface has been scratched

Future

- Need innovation with DB back-end
  - Still working with mysql issues
  - Still using RRD files
  - Need for scalable big-data collection
  - Moving into automation / SDN / OpenFlow

Questions

Thanks!

Aaron Hughes, Paolo Lucente,

Brent Van Dussen, Arien Vijn