

ITU-T Workshop on IP Traffic Flow Measurement

(Geneva, Switzerland, 24 March 2011)

Typical use of IP traffic flow measurements in telecom operators

**Paolo Lucente,
KPN International, pmacct**

Agenda

- **Introduction**
- **Traffic measurements:**
 - Who does what
 - Coupling flow data with BGP
- **Complications**

Internet: content to eyeballs a simplified taxonomy

- Players in the Internet:
 - A – Access providers (radio, cable, DSL, fiber, etc.)
 - B – IP carriers
 - B – IXPs, Internet eXchange Points
 - C – Content providers
 - C – Hosting

IP traffic flow measurements

- By this time of the day we should know lots of things already on the topic 😊
- Let's do a quick recap:
 - ➔ Flow-based technologies available: NetFlow, IPFIX
 - ➔ Efficient correlation of traffic data with routing information (BGP) is possible; tools are publicly available
 - ➔ Use of data reduction techniques (ie. sampling, aggregation) is valid to keep data-set manageable. Accuracy affected.
 - ➔ Use of divide-et-impera techniques is valid to distribute computing of results

Routing in the public Internet

- BGP is used for inter-domain routing in the public Internet
 - Traffic is routed to destination
 - Routing domains are distinguished by ASNs – Autonomous System Numbers:
 - Now 32-bit, “space for everybody”
 - Price for multi-homing is falling so people are even encouraged ...
 - Unlikely IP transit will disappear, things might well consolidate though ...
 - RPKI/ROA trend: traditional base of trust being broken
- => *“Can the basic block of inter-domain routing be mapped to something real, say, a country?” Not really!*

Agenda

- Introduction
- **Traffic measurements:**
 - **Who does what**
 - **Coupling flow data with BGP**
- Complications

Traffic measurements: common ground

- Capacity planning
 - ➔ Build capacity where needed
- Traffic Engineering
 - ➔ Steer traffic where capacity is available
- Security
 - ➔ Events notification, alarms, mitigation, etc.
- Historical traffic trends
 - ➔ Feed into Sales department
- Feed into 3rd party (even home-made) tools

Traffic measurements: content providers

- CDN (Content Delivery Networks):
 - ➔ Monitor traffic quality (mix of methods)
 - ➔ Destination-based accounting is popular:
 - Possible because one direction prevails ...
 - ... and this is the one routed by the CDN

=> *This is accomplished with flow-based IP traffic measurements*

=> *This is used more as an internal cost-control measure than to bill customers*
- Hosting, data-center:
 - ➔ Monitor co-located server quotas

Traffic measurements: IP carriers

- IP carriers use flow measurements to:
 - ➔ Detect revenue leaks
 - ➔ Determine customer profitability
 - ➔ Do customer retention
 - IP carriers interpretation of usage-based billing (typically SNMP-based):
 - ➔ Price per Mbps @ 95th percentile port utilization
 - ➔ Burst possible
- => *No distinction is made on IP traffic primitives when billing customers*

Traffic measurements: access providers

- Monitor customer quotas
 - Monitor BBA fair-usage policy compliancy
 - Abuse
 - Radius accounting is popular; but flow accounting is useful for abuse purposes
- => No distinction is made on IP traffic primitives when billing customers*

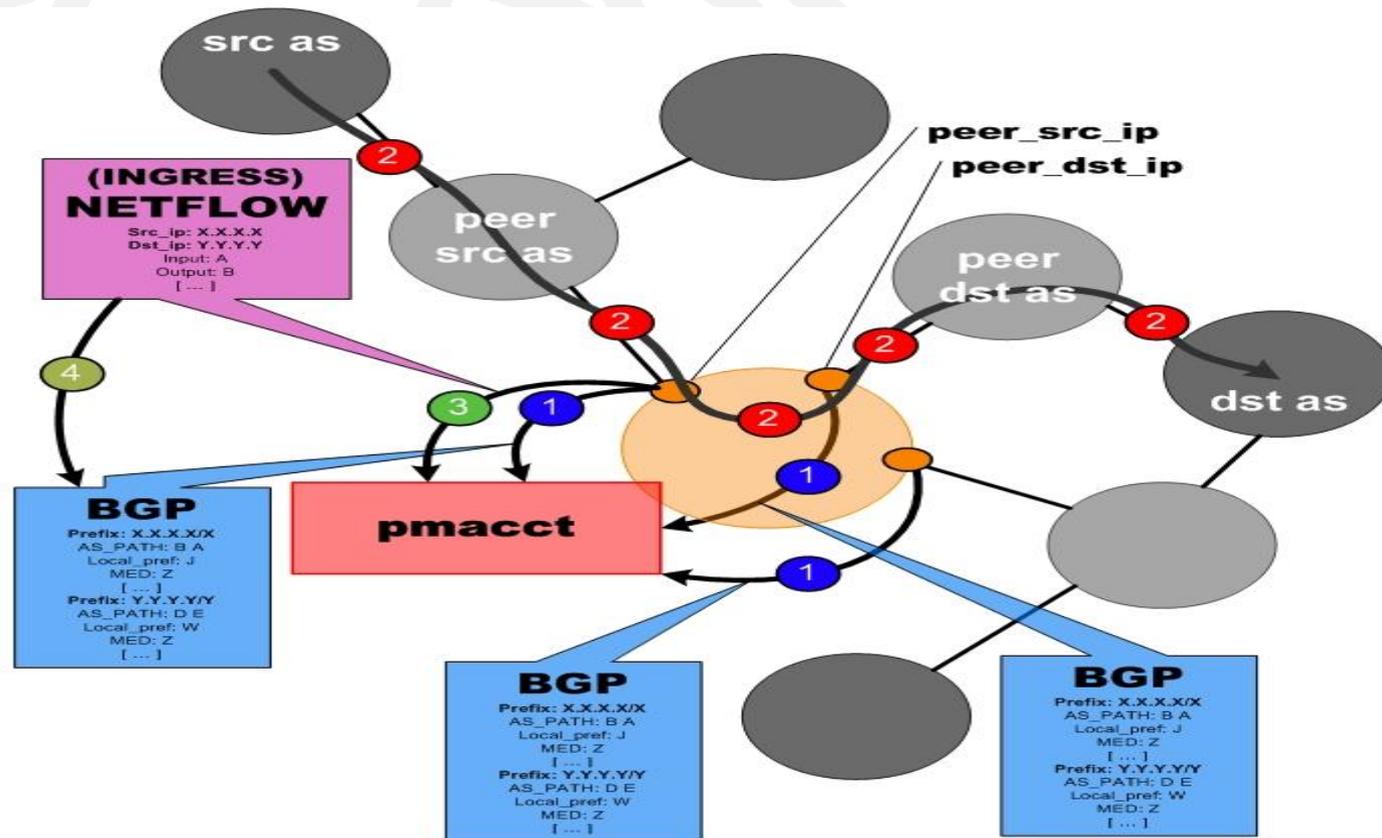
Traffic measurements: internal billing

- Networks with international scope
- Subsidiaries in several countries, responsible for own profits and losses
- Corporate international IP backbone
 - = > *Subsidiaries transit over the corporate IP backbone*
- Not uncommon backbone utilization is measured to bill subsidiaries proportionally
 - = > *This is accomplished with flow-based IP traffic measurements*

Coupling traffic measurements with BGP

- Method(s):
 - ➔ #1: Coupling at the router (feeling like moving control plane information over and over ..)
 - ➔ #2: Coupling at the collector (illustrated next slide)
- Tools are available for each method:
 - ➔ #1: ie. NetFlow v9/IPFIX, Cisco FNF, etc.
 - ➔ #2: ie. Arbor Peakflow (Commercial), pmacct (OSS)

Coupling traffic measurements with BGP (collector approach illustrated)



- 1 Edge routers send full BGP tables to pmacct
- 2 Traffic flows
- 3 NetFlow records are sent to pmacct
- 4 pmacct looks up BGP information: NF src addr == BGP src addr

Agenda

- Introduction
- Traffic measurements:
 - Who does what
 - Coupling flow data with BGP
- **Complications**

Control plane woes

- Control plane vs data plane
 - ➔ BGP says to route to X via Y and Z
 - ➔ Hidden more specifics in, say Y, can route to X via J

=> *Accuracy of accounting is jeopardized*
- Asymmetric routing, load-balancing and multi-homing impact needs a careful analysis

Application layer woes

- A traffic flow between two end-points, say A and B, is better represented by two uni-directional flows, say, A->B and B->A
 - In voice and TDM in general morphology of these two uni-directional flows is congruent
- => In the IP world this is never a guarantee. Morphology is dictated by the specific application.*

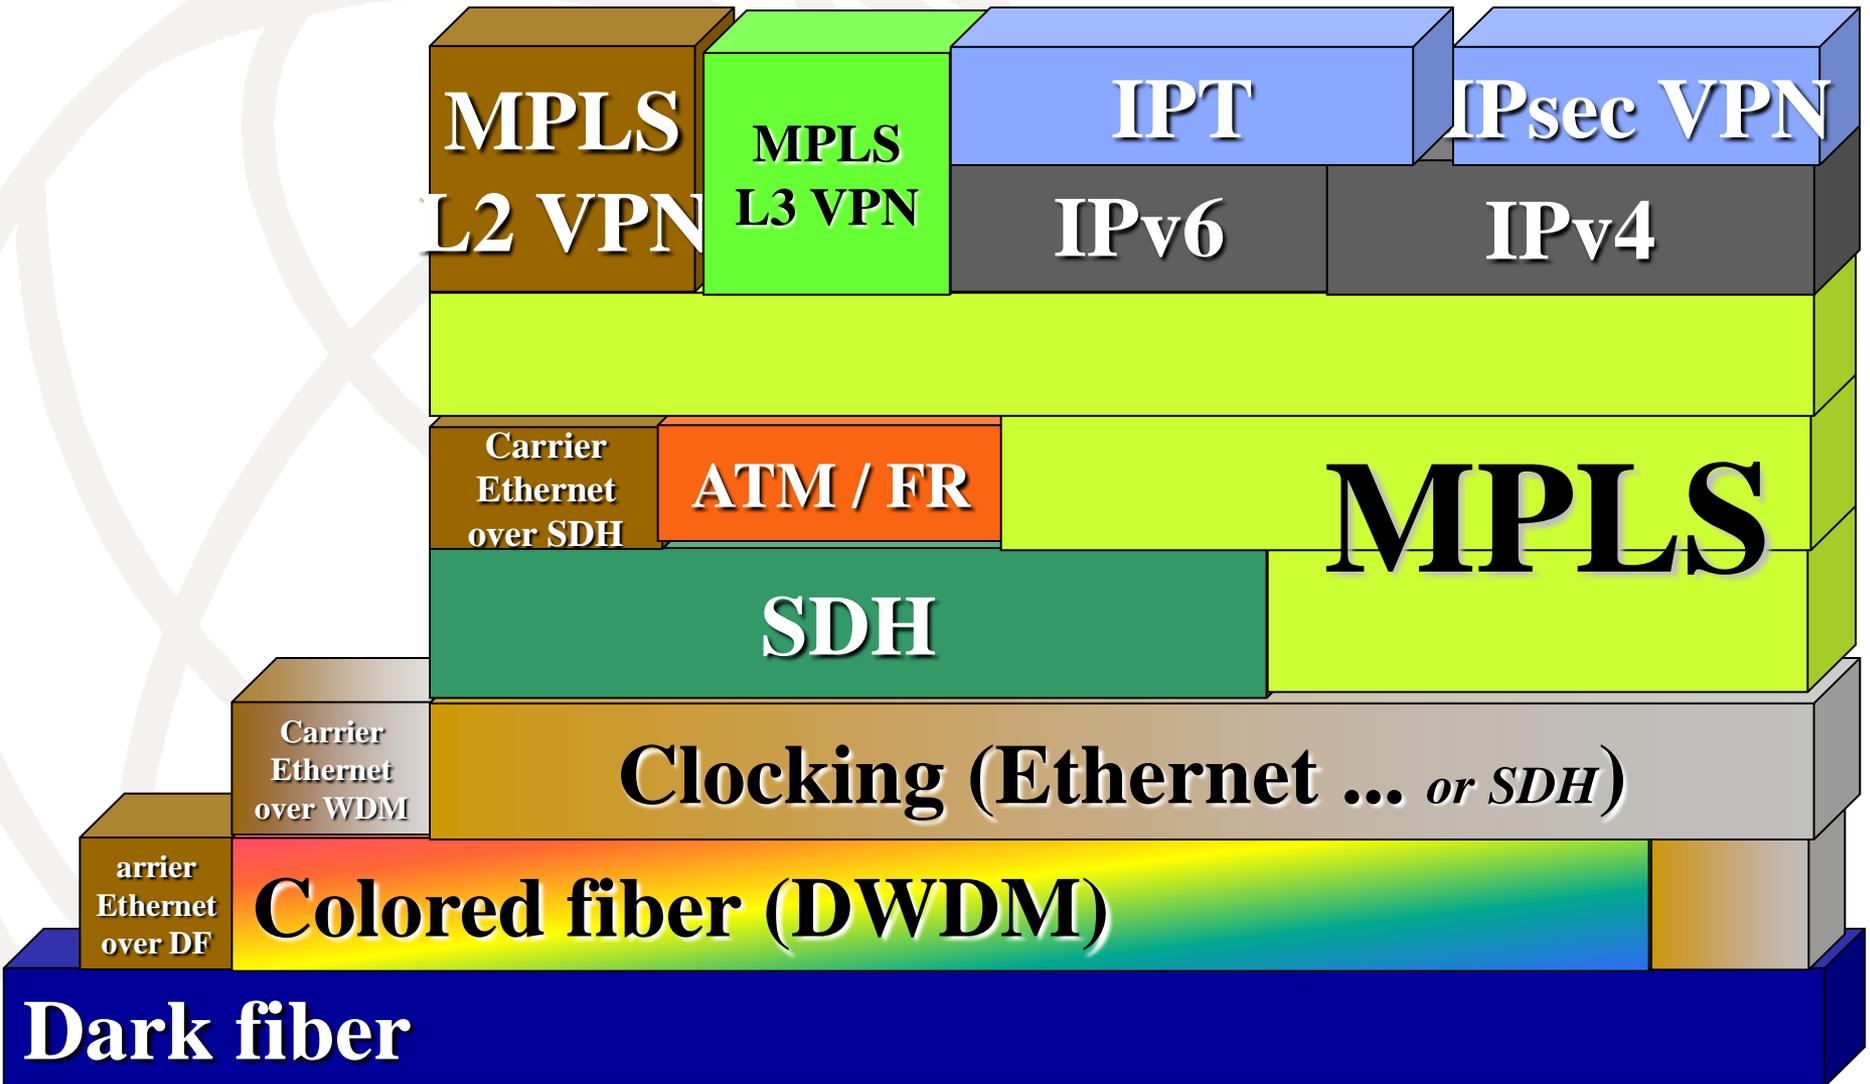
IP addressing woes

- IP addresses issued in one world region are not required to stay within that region
- Common case among wholesalers: they tend not to like to affiliate to multiple RIRs (costs, trainings, procedures, etc.)

=> *But:*

- *They lend part of their address space to customers*
- *They are not the best at documentation ☺*
- *Big dynamic IP pools worsen the situation*

Layering woes (L1, L2 and L2.5) (1/2)



Layering woes (L1, L2 and L2.5) (2/2)

- ISP A lies in country X
 - IP connectivity is cheaper bought in country Y
 - ISP A thinks of a combined solution:
 - ➔ They will buy transmission (dark fiber, wave, etc.) between X and Y from party B
 - ➔ They might optionally build a footprint in Y
 - ➔ They will buy IP transit in country Y from party C
- => *Which country this traffic belongs to?*

Thanks for your attention
Any questions?

Paolo Lucente

paolo.lucente@kpn.com

paolo@pmacct.net

Geneva, 24 March 2011