Advanced Traffic Analysis Techniques for Peering Networks, Utilizing Netflow.

Richard A Steenbergen <ras@nlayer.net> nLayer Communications, Inc.
Nathan Patrick <nathan@sonic.net> Sonic.net, Inc.

Why you need to know about your traffic

- To decide if you should peer with a new network.
- To convince other networks to peer with you.
- To manage traffic engineering to other networks.
- To defend your network against depeering actions.
- To make intelligent transit purchasing decisions.
  - Maximize your peering strategies.
  - Pick providers who are best for your specific traffic.

How to study your traffic? Netflow of course.

- Hopefully everyone has used or heard about Netflow, but just incase you've been in a coma:
  - Netflow is a simple framework for exporting summarized information about the packets being routed through your network.
  - Periodically this data is exported to a collection host via UDP.
  - External tools can parse these flow records for statistical analysis.

So what is wrong with existing Netflow?

- Netflow exports are good at telling you about the current state of the network.
  - Where packets are going now.
  - Some simple information about origin-AS or peer-AS.
- To be effective for peering strategy, you must expand on this information and become predictive.
- The ultimate question is not where **DO** you route your traffic, it is where **CAN** you route your traffic.

Ok already, tell us the new techniques

- Start by throwing out (almost) all information from the flow export except the destination address and the total octet count.
- Build your own virtual RIB(s) using externally collected routing information.
  - Prefixes and AS-PATHs from a given point of view.
- Almost all further analysis is just a matter of changing the RIBs or the AS-PATH position.

A word about why this works: Multihoming.

- Multihoming is pervasive at the core. Even if you don't multihome, your Tier 2 transit provider probably does.
- Empirical evidence suggests that the average Tier 1 has less than 10% of its customer base single-homed.
  - Or: 90% of the customers you can reach through someone else.
- BGP obscures alternate paths with every hop and every best-path decision. Once this data is gone, there is no way to get it back.
  - The only solution is to look at routes from different views.

Application: Predicting traffic to a new peer.

- Collect the peer's customer routes via OOB.
  - 111.2.0.0/16 1234 7183 7164 2616 143
  - 111.3.0.0/16 1234 7183 3834 818 82
  - 111.80.0.0/17 1234 829 817 646 7173
  - etc
- Set *n* = 1 (examine the first AS in the PATH)
- Project traffic onto this RIB, counting bits that would hit the AS at position *n*.
- You now know about your total traffic to ALL of a potential peer's customer routes.
- You can expand on this by examining the Netflow nexthop or Peer AS to determine where you send the traffic today.

Application: The art of persuasive peering.

- Some networks are aggressively open peering ("Peerleaders"), other networks take a little convincing.
- Often times, they just don't have the right data.
  - Billion dollar networks aren't necessarily any better off when it comes to understanding their traffic.
  - Inbound traffic is much harder to predict than outbound. The outbound network may have insights that the receiver of the traffic simply doesn't.
  - Who needs hard data when you have ideology and company Kool-Aid?
- Having "proof" to back up your claims is a good way to get noticed out of a crowd of folks with Linux routers and a "Global" "Fully Redundant" "OC-192" 0-Commit $500 MPLS "Backbone".

Application: Donut Peering

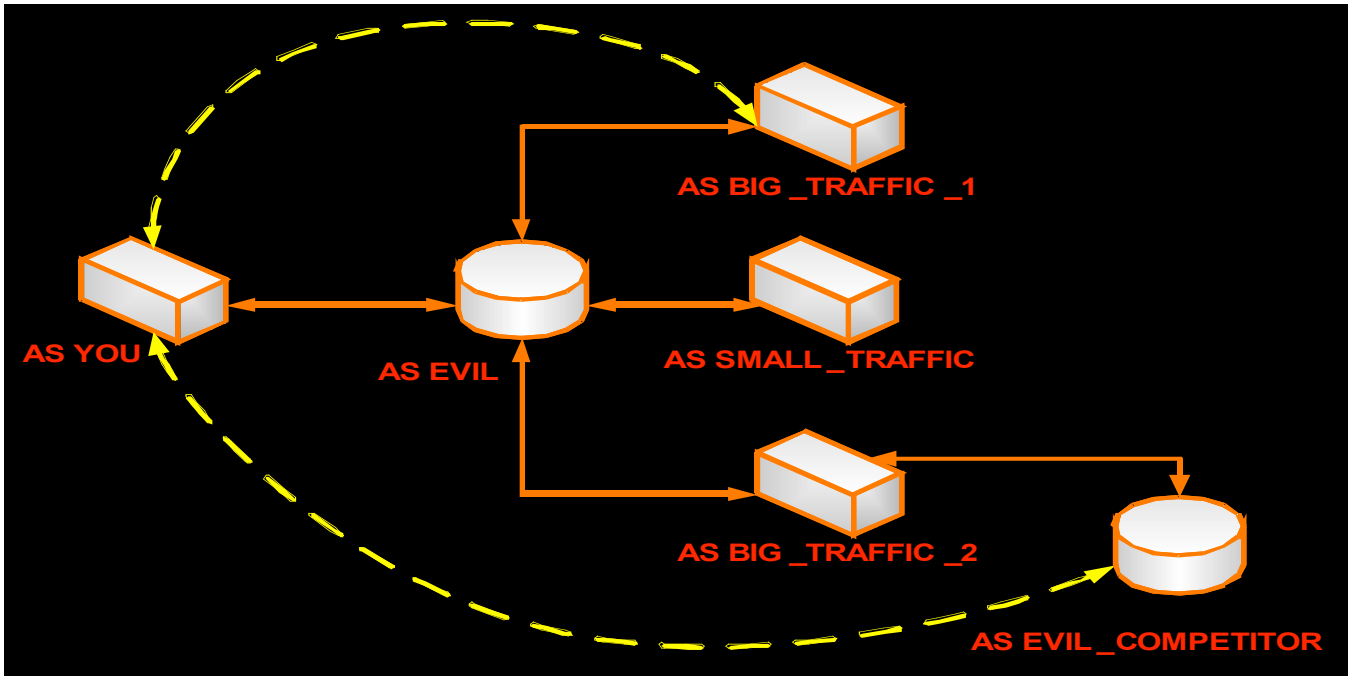

- Some networks just won't peer with you, no matter how much technical or financial sense it makes.
- If you can't work with them, try working around them.

Application: Donut Peering

- As always, you have several options:
  - Try to peer with their customers.
  - Try to sell to their customers.
  - Try to find their customers' customers.
- Obtain a RIB for the Peer in question:
  - *n* = 1 yields total traffic.
  - *n* = 2 yields traffic to their specific customers.
  - If necessary, obtain a RIB for the specific customers. Remember, Customer may have more routes!

Application: Picking your Transit Providers

- How do you pick your transit providers? A good price and a smooth sales pitch, or based on hard data?
- The same analysis works on a provider's RIB too:
  - By understanding where a particular transit provider sends your traffic, you can better understand their routing policies and which networks may need special attention.
  - Try our new transit providers virtually, before you buy.
  - Pick transit providers who support your peering strategy. It may make sense to buy transit from someone who doesn't already send traffic to your potential peers.

Examples: Sprint (AS1239) (or: Show me some pretty pictures already)

- Just how much can an average network Donut?
- Let's look at this graph showing traffic to Sprint:

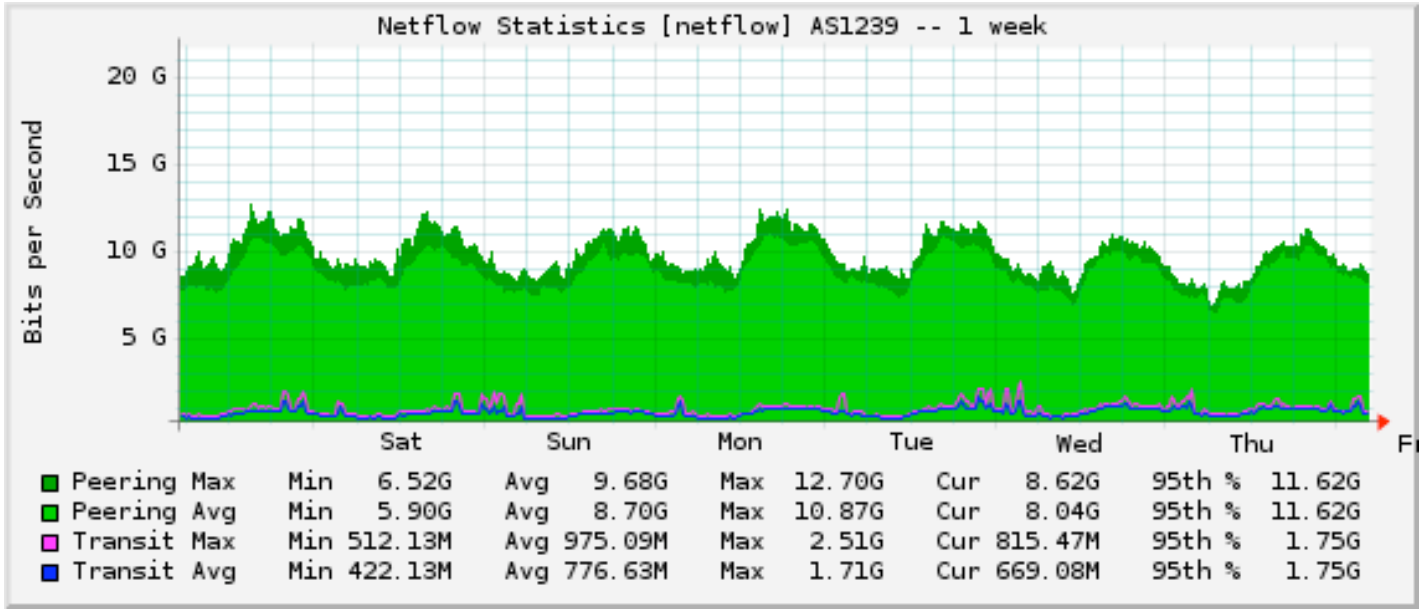

Examples: TeliaSonera (AS1299)

- Thanks to Peter Cohen for being a willing victim.
- A simple traffic graph from a medium-sized NSP:

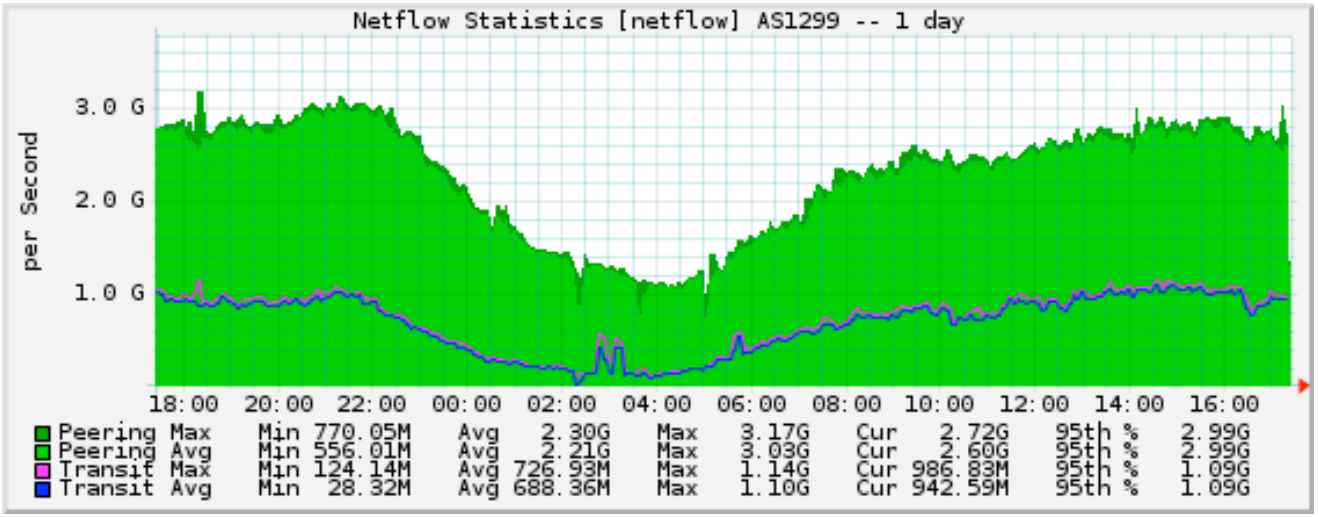

- What this says:
  - Out of all of the customer routes of AS1299, this network already peers out 3 Gbps, but sends 1 Gbps to their transit(s).

Examples: TeliaSonera (AS1299)

- An analysis of where they send those 4 Gbps:

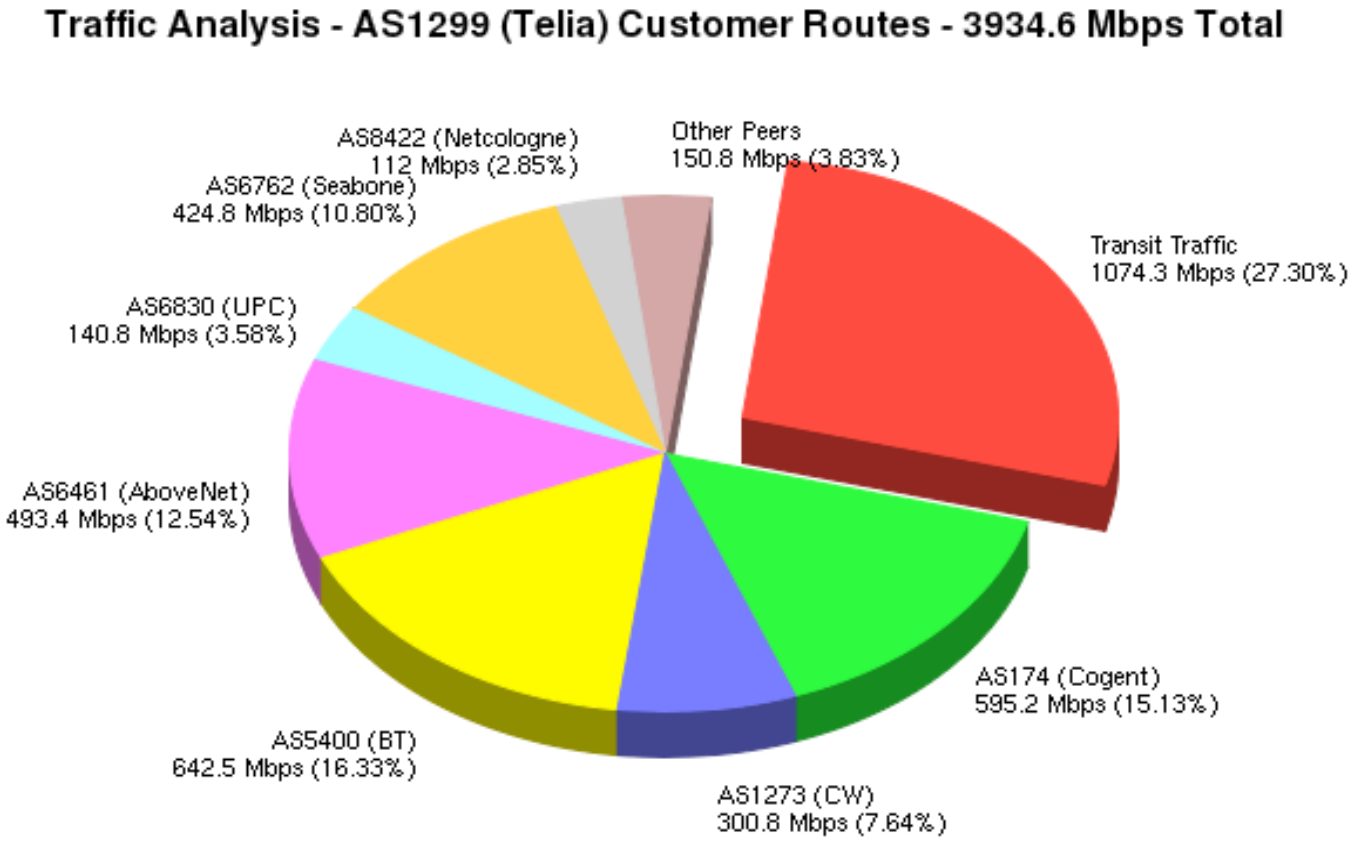

Examples: TeliaSonera (AS1299)

- An analysis of where their transit providers send that previously mentioned 1 Gbps:

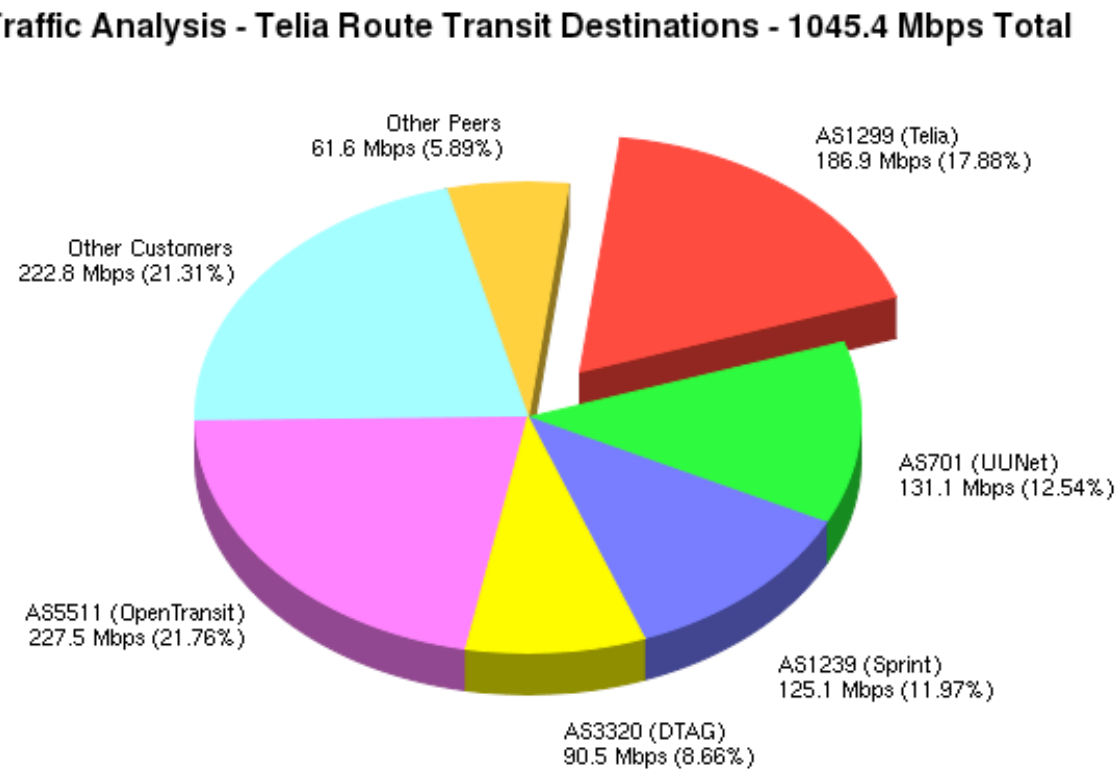

Conclusion: TeliaSonera (AS1299)

- AS1299 carries only 187 Mbps (or 4.88%) of the potential 4 Gbps of traffic sent by the example network.
- The rest of the traffic bypasses them completely
  - Goes directly to their multihomed customers, or
  - Worse still, goes to their competitors.
  - Either way, this is traffic they will never be able to bill for.

- By looking at the next AS hop, we have a list of their customers, and how much traffic is sent to each.
  - Convincing: Telia can calculate additional revenue from peering.
  - Peering/Poaching: You now have a list of the customers you send the most traffic to. If you can peer around them, Telia may become irrelevant to you.

Flaws in the system (or: You knew it wasn't going to be this easy!)

- So far we've only talked about outbound traffic
  - That's because inbound is far more difficult to predict.
  - Remember that the outbound network is in complete control, and your inbound is someone else's outbound.
- Gathering RIBs is hard work.
  - No existing route-servers collect "peer views".
  - Many networks consider this proprietary information.
  - A large percentage of the data can come from public looking glasses.
- Traffic will shift as AS-PATH lengths change.
- You won't accept every prefix of a potential peer, and simulated best path calculations are too difficult to predict in a complex network.

Ok now give me a tool that does this stuff

- http://asflow.sourceforge.net
- A simple tool for text-only version, available in two flavors:
  - Perl
    - Pros: Incredibly simple, uses existing flowtools data captures.
    - Cons: Slow and consumes a lot of memory. Intended for quick use against existing “5 minute sample” captures.
  - C
    - Pros: Much better memory usage and run-time CPU usage.
    - Cons: Much more complex, designed for long-term use.

Other resources

- Packet Clearing House peer views for RIBs
  - http://lg.pch.net
  - http://www.pch.net/resources/data/routing-tables/archive/
- Other looking glass views
  - http://www.traceroute.org
  - http://www.bgp4.net

Send questions, complaints, threats, etc. to:

Richard A Steenbergen <ras@nlayer.net>

Nathan Patrick <nathan@sonic.net>