S-38.3192 Verkkopalvelujen tuotanto
S-38.3192 Network Service Provisioning
Lecture 9: Peering

Part of the material presented in these slides is based on BGP lectures of Olivier Bonaventure www.info.ucl.ac.be/people/OBO/BGP/

Internet

- **The value of Internet is in global reachability**
  - Reachability comes from co-operative peering efforts
    - Customer peering (Customer-Provider-Customer relationship)
    - Shared cost peering (Provider-Provider relationship)
- **There are roughly 48000 assigned AS numbers from which about 24000 are advertised to BGP i.e. are players**
  - 21165 of them are Stub ASs
  - 83 are pure transit providers
  - 3082 do both

Internet

- The structure of Internet is chaos
  - Thousands of service providers with highly varying principles in their operation

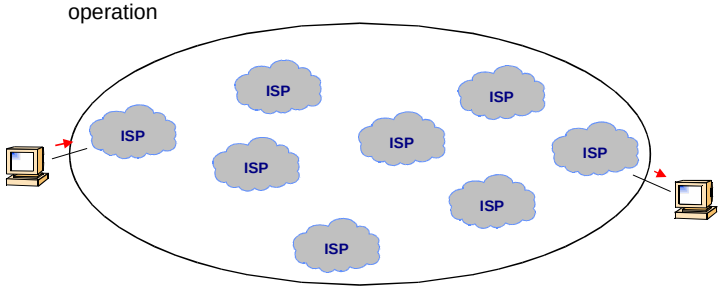

Internet

- **How packet finds its route through the black box**
  - BGP forms a structured layout of the whole Internet for packet level transport
    - Reflects the semi-optimal contractual agreements between operators along the route of the packet
- **Why accepting packets from fellow ISP**
  - Economic impact
    - Transit traffic
  - Reciprocity
    - Cost reduction

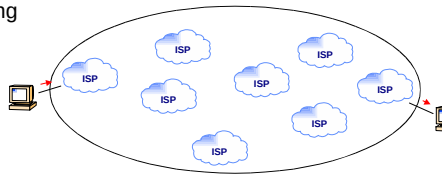

Agreements

- Form the basis between inter-provider communications
  - Small ISPs are customers of larger ones
    - Larger ISPs deliver their customer traffic as their own traffic
    - Larger ISPs deliver their customer traffic as transit traffic
  - Equal size providers exchange their traffic pro bonus
    - Both save money by interconnecting directly rather than through 3rd party
    - Mutual agreement for exchanging only their customer traffic

Strict hierarchy

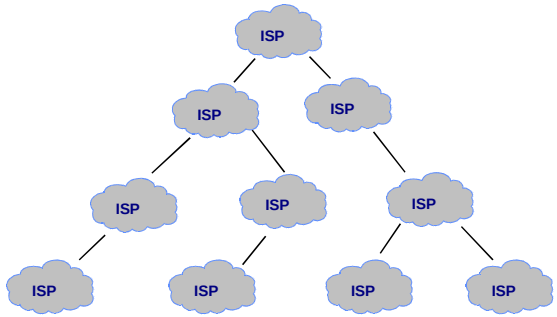

Strict hierarchy

- Based on structural and regulated manner of forming customer/provider relationships
  - Valid in telco operations
  - Operators for a chain of customer/provider relationships
  - Based on regulation of operational arena
    - Local operators
    - Long distance operators
    - International operators
  - Cash flows to the top of the hierarchy
    - Local operators collect the money from end users
    - Middle layers take their premiums

Loose hierarchy

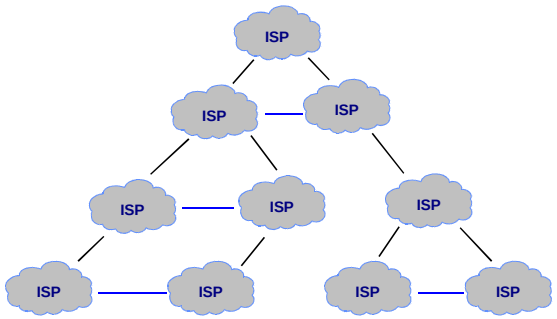

Loose hierarchy

- Local providers compete the local market but share common need to exchange their customer traffic on a local level
  - It is profitable for all to have direct exchange of traffic without 3rd parties
    - Better marginal revenue
  - Requires
    - Interconnection points
    - Bilateral agreement to establish equality
      - Zero payment principle
        - » Both parties benefit from peering
        - » No mutual transfer of money

Internet

- Naturally loose in hierarchy
- Local ISPs maximize their revenue by minimazing their transit traffic
- Same structure on all levels of hierarchy
- Any connection through the Internet is formed with chain of customer/provider relationships with a single zero payment border
  - Cost of connection is therefore divided into two
    - From source to top of the chain
    - From destination to top of the chain
  - Peering does not cover transit traffic
    - Only one zero payment border

Internet

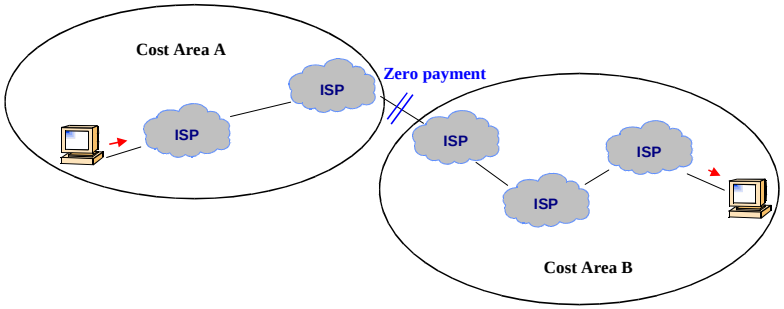

Transit domain

- A transit domain allows external domains to use its own infrastructure to send packets to other domains

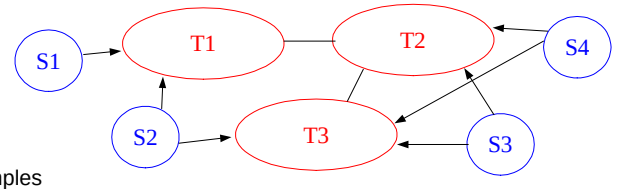

- Examples
  - FuNET, NorduNET, GEANT, Internet2, BT, Telia, Level3,...

Stub domain

- A stub domain does not allow external domains to use its infrastructure to send packets to other domains
  - A stub is connected to at least one transit domain
    - Single-homed stub : connected to one transit domain (S1)
    - Dual-homed stub : connected to two transit domains (S2-S4)

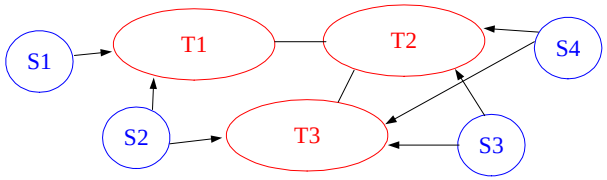

Stub domain

- Examples:
  - Content-rich stub domain
    - Large web servers : Yahoo, Google, MSN, TF1, BBC,...
  - Access-rich stub domain
    - ISPs providing Internet access via CATV, ADSL, ...
      - Saunalahti, Kolumbus, Welho etc

Internet

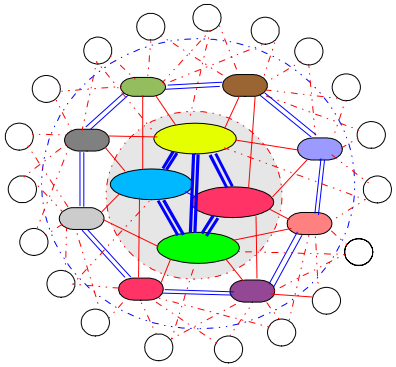

- Tier-1 ISPs
  - Dozen of large ISPs interconnected by shared-cost peering arrangements
  - Form the core of the Internet
  - Provide transit service for T2/T3 service providers

Tier-1 service providers

- AT&T (AS7018)
- Global Crossing (GX) (AS3549)
- Level 3 (AS3356)
- Verizon Business (formerly UUNET) (AS701)
- NTT Communications / (formerly Verio) (AS2914)
- Qwest (AS209)
- SAVVIS (AS3561)
- Sprint Nextel Corporation (AS1239)

- Dozens of other providers are also marketed as T1 where as they are not
  - They have
    - Transits
    - Paid peering

Internet

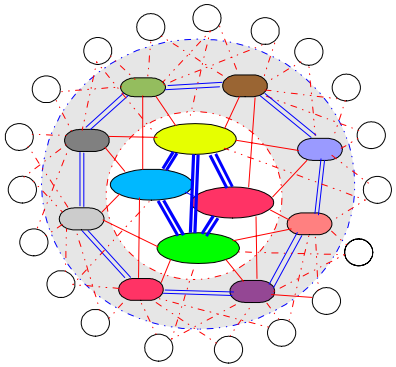

- Tier-2 ISPs
  - Regional or National ISPs
  - Customer of T1 ISP(s)
  - Provider of T3 ISP(s)
  - shared-cost with other T2 ISPs

Internet

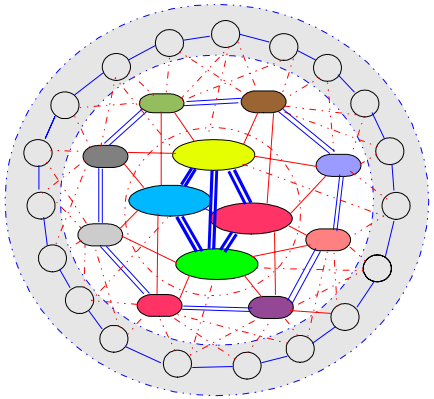

- Tier-3 ISPs
  - Smaller ISPs, Corporate Networks, Content providers
  - Customers of T2 or T1 ISPs
  - shared-cost with other T3 ISPs

Customer-provider peering

- **Principle**
  - Customer sends to its provider its internal routes and the routes learned from its own customers
    - Provider will advertise those routes to the entire Internet to allow anyone to reach the Customer
  - Provider sends to its customers all known routes
    - Customer will be able to reach anyone on the Internet

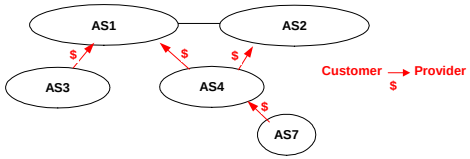

Shared-cost peering

- **Principle**
  - PeerX sends to PeerY its internal routes and the routes learned from its own customers
    - PeerY will use shared link to reach PeerX and PeerX's customers
    - PeerX's providers are not reachable via the shared link
  - PeerY sends to PeerX its internal routes and the routes learned from its own customers
    - PeerX will use shared link to reach PeerY and PeerY's customers
    - PeerY's providers are not reachable via the shared link

Shared-cost peering

- AS1 send routes of AS{1,3,4,7} to AS2
- AS2 sends routes of AS{2,4,7} to AS1
  - Not AS3 while those routes come from shared-cost peering
    - Routes from shared-cost peering are not advertised to providers

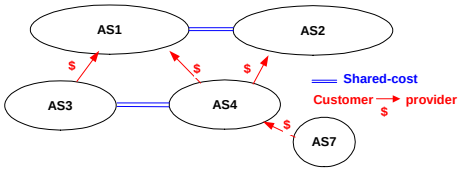

Direct Approach

- In direct approach peering negotiations are established with open cards
  - Official invitation to peered AS to start negotiations
    - Results
      - Peering formed
        - » Direct
        - » Partial
        - » Conditional
      - Peering not formed

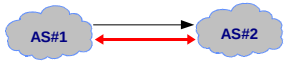

Migration tactics

- Initially transit connection is bought with option to peer when conditions are met
  - Incentive of accumulated cash flow for period of time
  - Risk of not having peering conditions met or changed conditions over the time

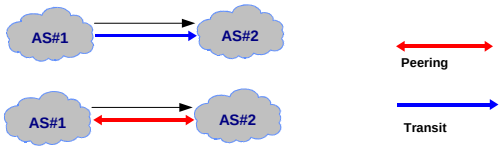

Selected transit

- Transit connection is bought from provider which is not a candidate of future peering

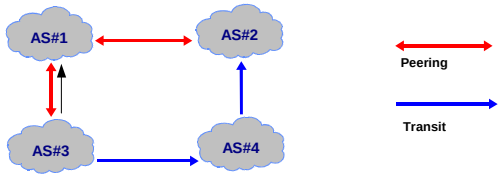

Traffic Manipulation

- Force traffic to routes that make peering look attractive
  - After while initiate peering negotiations
    - Upstream traffic does not generate fast incentives

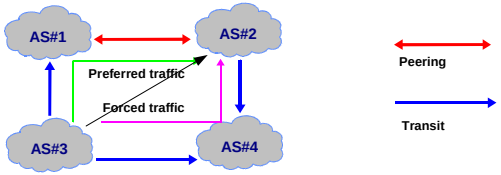

Traffic Manipulation

- Force traffic to routes that make peering look attractive
  - Stop advertising routes to cheap directions
  - Falsely inject AS#2 to path vectors in direction of AS#1
    - Loop detection prevents routes to be installed at AS#2

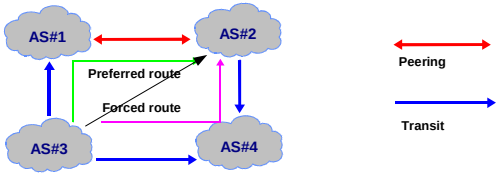

Internet

- **Local providers aim to minimize their expenses by interconnecting at local level**
  - **Local exchange points**
    - ..CIX (Commercial Internet eXchange)
    - MAE.. (Metropolitan Area Exchange)
    - NAP (Network Access Point)
    - IXP (Internet eXchange Point)
    - EP (Exchange Point)
  - **Bilateral interconnections**

Internet exchange

- **Commercial starting point**
  - A company builds an interconnection point to
    - Gain revenue from peering traffic
    - Gain revenue from transmission links coming to exchange
    - Gain revenue from transit traffic
- **Co-operative starting point**
  - Neutral partner runs the exchange
    - None of the partners owns the premises
    - None of the partners owns the transmission links into exchange
    - None of the partners owns the equipment in exchange

Internet exhange

- Build over L2 technology
  - Ethernet, ATM, FrameRelay switch
- Each provider connects into shared media with transmission link terminated to border router of provider
  - Everybody is able to see everybody

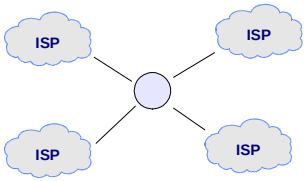

Internet exhange

- Peering agreements can be based on
  - **Multilateral agreements**
    - Every partner is peering with every other partner
      - All border routers share a common subnet which is not filtered
        - » Ideal situation for Ethernet type of IXP solution

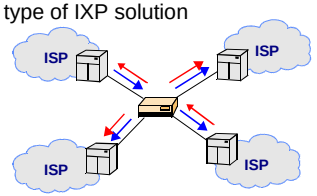

Internet exhange

- **Bilateral agreements**
  - Partners peer only based on bilateral agreements
    - Requires L2 technology that is able to create virtual connections between peering partners
      - » ATM PVC
      - » FR DLCI
      - » Ethernet VLAN

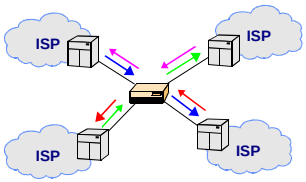

Internet exhange

- **Multilateral peering reguires either**
  - **Separate BGP session between each border router**
    - **N(N-1) sessions**
  - **IXP offers route server capabilities**
    - Only N sessions
      - BGP-route reflector

Internet exhange

- Depending on operational philosphy of IXP
  - **Partners can make bilateral transit agreements in IXP**
    - Partners are already in same premises
    - Required separate virtual connections between transit provider and customer
  - **Partners can make QoS peering**
    - Several virtual connections between peers
      - One per VPN per QoS class
      - One per MPLS LSP
      - etc

Material about peering

- Peering
  - As general
    - http://en.wikipedia.org/wiki/Peering
  - Motivations
    - http://www.equinix.com/pdf/whitepapers/PeeringWP.2.pdf
  - Tactics
    - http://www.nanog.org/papers/playbook.doc
- IXP
  - How-to
    - http://www.euro-ix.net/ixp/startingixp/