Internet Routing Protocols
Lecture 05
Loc/ID split to the Rescue?

Advanced Systems Topics

Lent Term, 2010

Timothy G. Griffin
Computer Lab
Cambridge UK

Inter-domain routing is often asymmetric

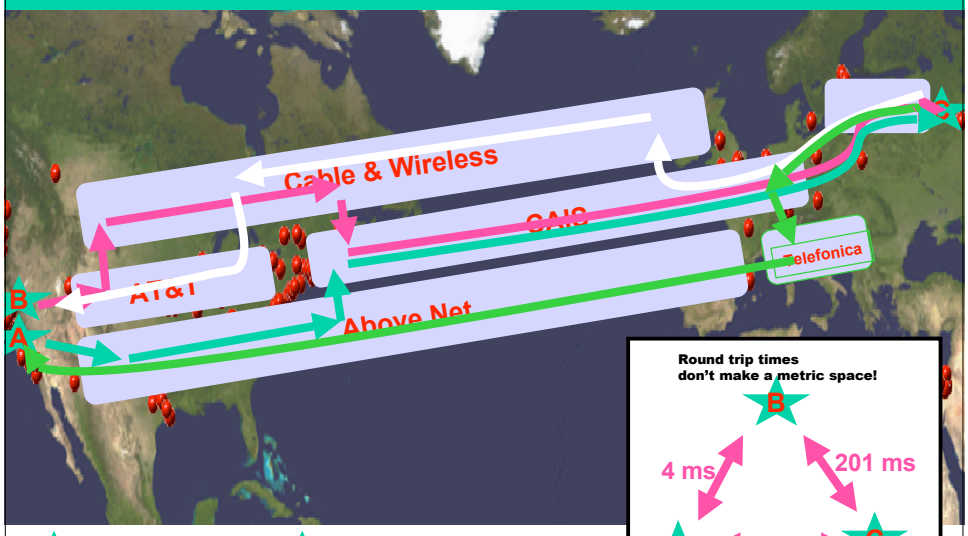

A = HP Palo Alto
C =Moscow State U
B =Intel Berkeley

What is causing BGP Routing Table Growth?

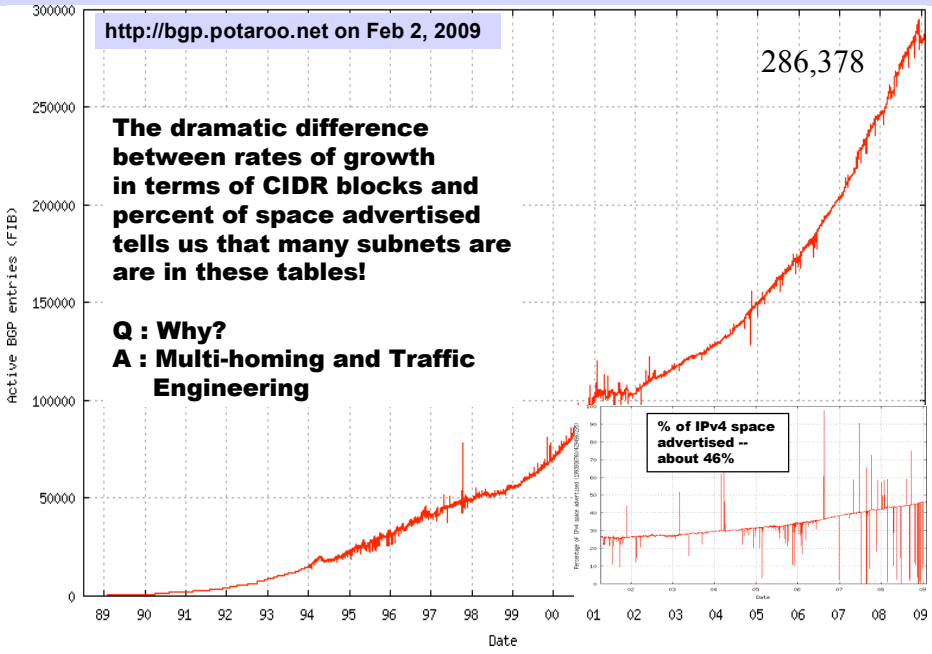

Deaggregation Due to Multi-homing

If AT&T does not announce the more specific prefix, then traffic to MODUS MEDIA will go through Genuity because it has a longer match....

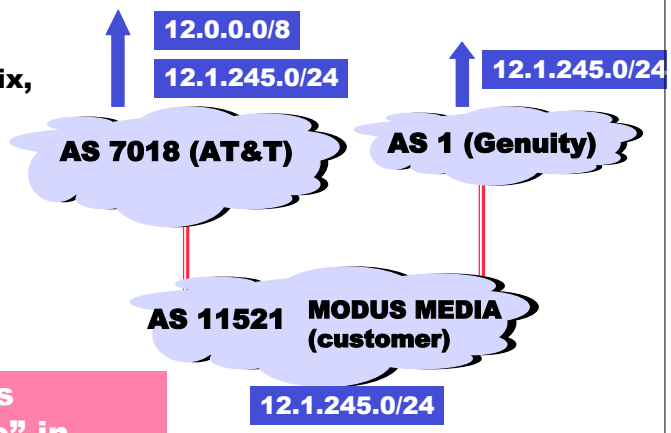

MODUS MEDIA is "punching a hole" in the 12.0.0.0/8 CIDR block

Deaggregation Due to "Traffic Engineering"

Remember: control of inbound traffic (with outbound routes) is very difficult, so network operators use whatever hacks they can get their hands on!

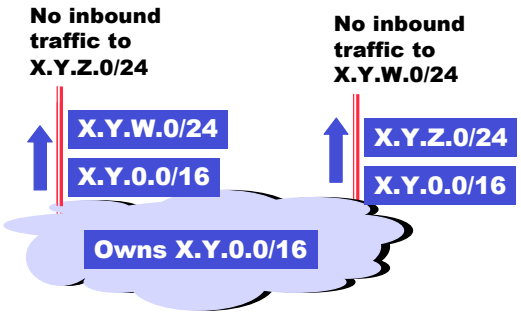

Loc/ID Split as an Architectural Solution?

- Problem: IPv4 and IPv6 addresses have overloaded semantics
- Conceptually, we have two distinct address spaces
  - Endpoint IDs (EIDs) --- public IP address space
  - Routing Locators (RLOCs) --- infrastructure (backbone routers, links)
- These are conflated today, and EIDs aggregation is failing since it is not congruent with infrastructure topology
- Basic idea of Loc/ID split :
  - Packet to EID destination d hits an Ingress Tunnel Router (ITR) in backbone
  - The ITR finds a mapping (somehow!) of EID d to Locator l
  - The ITR encapsulates packet, sends to l
  - Encapsulated packet reaches Egress Tunnel Router (ETR) at l, which strips off encapsulation and sends traffic on to d
- A Loc/ID split would allow
  - topological addressing for Locators
  - Much smaller routing tables in the backbone
  - More control over inbound traffic (via the mapping function)
- But, would require
  - Control plane: A new means of mapping EIDs to Locators
  - Data plane : Encapsulation in the backbone

Loc/ID split

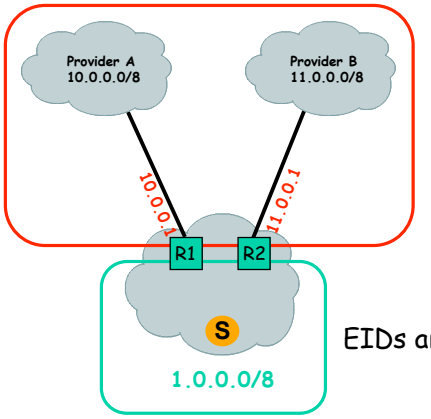

RLOCs used in the core

Mapping Database Entry:

1.0.0.0/8 -> (10.0.0.1, 11.0.0.1)

EIDs are inside of sites

Use Map-n-Encap

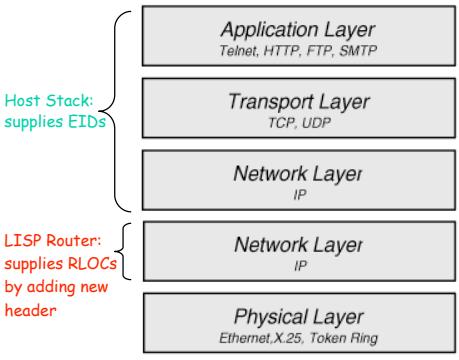

Mapping Entry:

EID-prefix: 2.0.0.0/8

Locator-set (RLOCs):

12.0.0.2, priority: 1, weight: 50

13.0.0.2, priority: 1, weight: 50

The universal solution to all problems in CS : introduce a layer of indirection!

Unicast Packet Forwarding

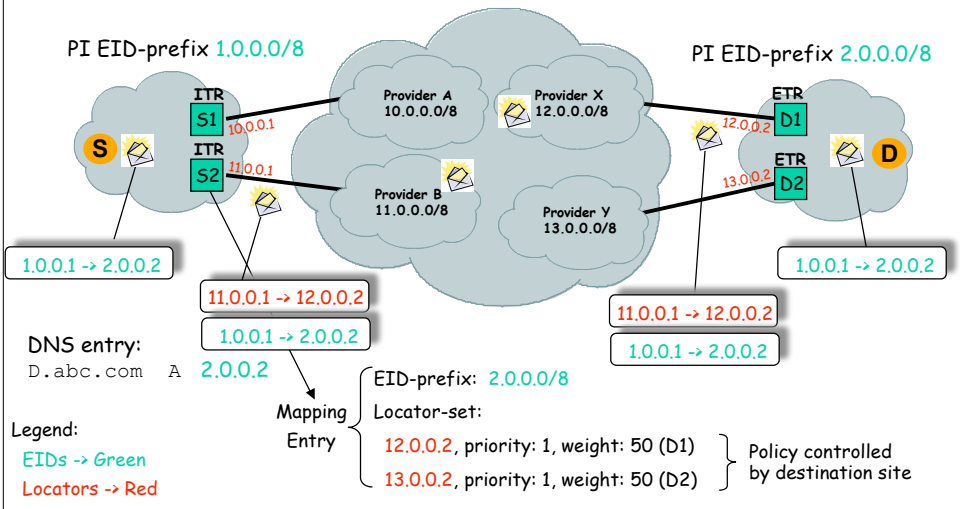

Many problems, many solutions

- A mapping service needs to be implemented
  - Current front-runner is called BGP-ALT
  - Idea : Run two instances of BGP
  - One BGP instance runs on real topology of locators
  - Another (alternate universe) instance of BGP runs on a virtual (overlay) network constructed with tunnels.
    - Assumption (untested, but reasonable) : since this is not tied to real topology, the EID space can be highly aggregated
    - Virtual network is used only for sending mapping requests to mapping servers
- A protocol is needed to communicate mapping info (request/reply)
  - Current front-runner is the Locator Identifier Split Protocol (LISP)
    - Network-based solution
    - No changes to hosts whatsoever
    - No new addressing changes to site devices
    - Very few configuration file changes
    - incrementally deployable
    - Address family agnostic
- Transition
  - Too complicated to get into!

LISP+ALT Control Plane

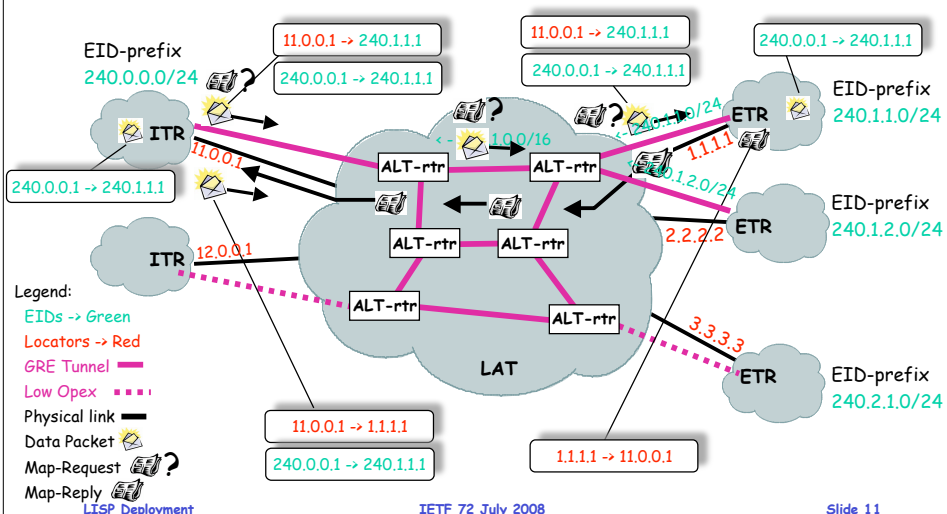

Other Benefits

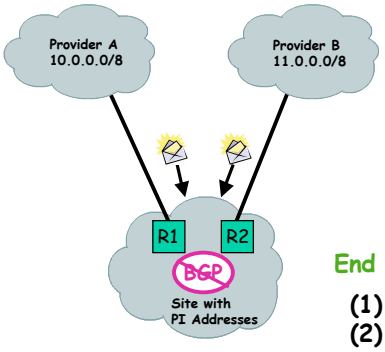

Provider Benefit

(1) Improve site multi-homing
(2) Improve provider traffic engineering
(3) Reduce size of core routing tables
(4) This implies much less update "churn"

End Site Benefit

(1) Easier Transition to IPv6
(2) Change provider without address change
(3) Better control of inbound traffic

For more Loc/Id split info

- Routing Research Group (RRG)
  - http://tools.ietf.org/group/irtf/trac/wiki/RoutingResearchGroup
- LISP Internetworking
  - http://www.lisp4.net/

  Some slides of this lecture were lifted from this site