

Scaling the Global Routing System: A Quick Overview of Ongoing Efforts

Lixia Zhang

November 7, 2008

Outline


- ◆ What problems facing us today
- ◆ Why we have these problems
- ◆ How to solve these problems (at high level)

Disclaimer

- ◆ Wearing no hat
- ◆ Personal opinion

How Internet topology looks like today


- ◆ When we draw network graphs, it tends to look like this



How Internet topology looks like today

- ◆ When we draw network graphs, it tends to look like this

Over time lots of growth at edges

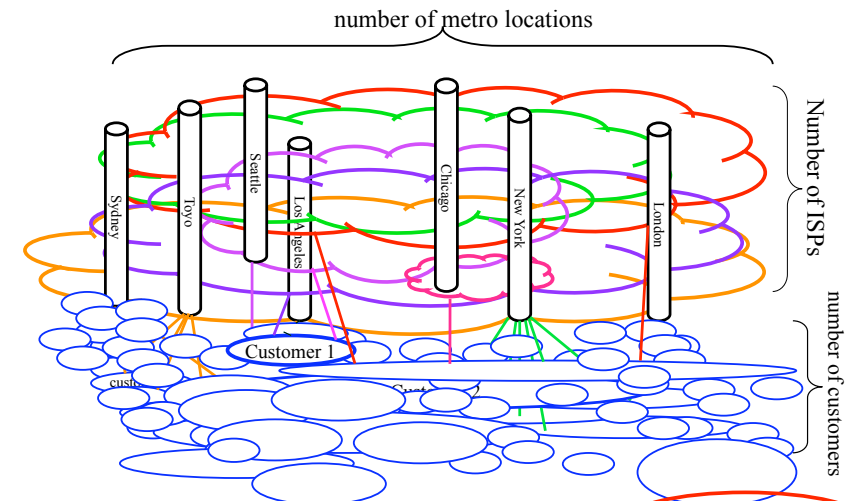


10/1/08

Gottingen

5

But in reality, it is more like this



DFZ Routing table size = Function(# of ISPs X # of PoPs X # of user sites X TE)

10/1/08

Gottingen

6

Do We Have A Problem to Solve?

- ◆ “Now the vendors are producing or planning to produce greater capability and higher performance routers. Is it possible to produce routers with higher performance to meet the rapid routing table growth?”

Do We Have A Problem to Solve?

- ◆ No: current router technology can support a routing table size at least 10 times bigger
 - And with Moore’s Law, this number will continue to go up over time
- ◆ Yes: the global routing table size facing *uncontrolled* growth at super-linear speed
- ◆ Facts:
 - Current routing table size does not reflect demand
 - No glorious record in predicting the future
 - It is not just availability, it is affordability
 - Network economics
 - Power consumption: the ultimate limitation

10/1/08

Gottingen

7

10/1/08

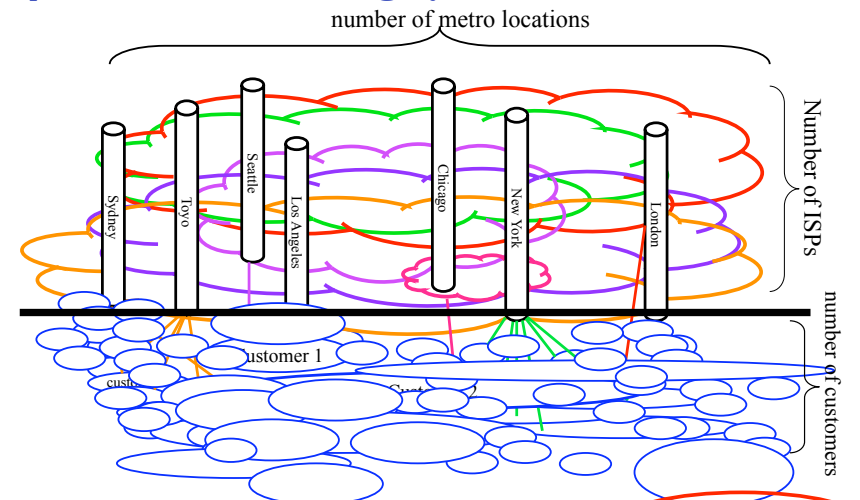
Gottingen

8

Defining Scalable Routing

- ◆ Being able to control the scale of the routing system
 - The ability to control, rather than any specific numbers
- ◆ Allowing the global transit core to route on *aggregatable* prefixes only

Proposed solution: Separate non-aggregatable prefixes from routing system in the transit core



DFZ Routing table size = ~~Function(# of ISPs X # of PoPs X # of user sites X TE)~~

Scalable Routing: Solution Space

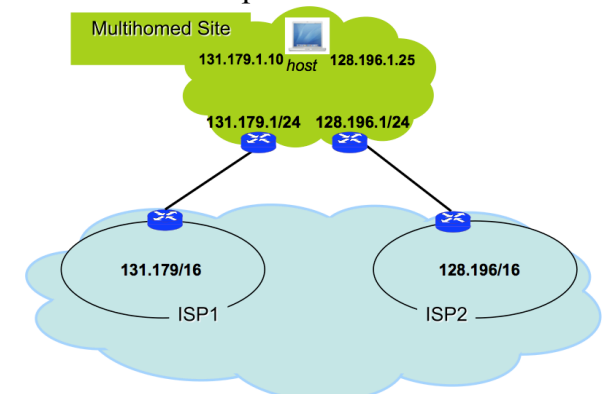
Two ways to get there

- ◆ **Elimination:** eliminating non-aggregatable PI prefixes from the *entire Internet*
- ◆ **Separation:** separating (removing) non-aggregatable PI prefixes from the global routing system

"Towards A Future Internet Architecture: Arguments for Separating Edges from Transit Core", Dan Jen et al.
ACM Workshop on Hot Topics in Networks (HotNets), October 2008

Elimination

- ◆ Eliminating non-aggregatable prefixes from the entire Internet
 - multihomed sites take in multiple PA addresses



Proposals under Elimination

- ◆ SHIM6
- ◆ (ILNP)
- ◆ (Multipath transport, i.e. Trilogy/Handley's ongoing effort)
- ◆ Common issues for all:
 - Require end host changes *at all edge networks*
 - Require site renumbering whenever changing providers
 - Require a solution to enable traffic engineering by operators

Separating edge prefixes from transit core

- ◆ Map & Encap: A number of proposed solutions
 - APT, IVIP, LISP, TRRP
- ◆ Requires a mapping system to glue the edges through the middle

