

Scaling Inter-domain Internet Traffic Engineering

Dana Blair – Cisco Systems (dblair@cisco.com)

The number of customers with 2 or more connections to the Internet (Internet Multi-homing) is growing and will continue to grow. These customers need N+1 connections for redundancy for their business critical applications. Once redundancy is in place, these same customers quickly notice that BGP or static routing to the Internet do not take full advantage of all their ingress and egress paths or react to performance issues. Then they look for Internet Traffic Engineering tools to optimize the usage of their Internet connections, both egress and ingress, for performance, utilization, and in some cases cost. They may also subdivide their public address space to further enhance the Internet Traffic Engineering capabilities. Some customers/providers may develop their own custom tools.

The dramatic growth in Internet Multi-homing is inevitable. How will it scale ?

Questions:

1. Is BGP multi-homing by Internet Customers contributing to the increased growth rate of the Internet Routing Table and rate of BGP update messages ?
2. Is BGP multi-homing “good enough” for ingress Inter-domain Traffic Engineering ?
3. Egress Internet Traffic Engineering can cause rapid changes in traffic patterns inside the Internet ? What’s the impact on interconnects between Internet Core providers ?
4. What is the correlation between Routing Instability (lack of convergence) and Application Performance ? Which route instability events do or do not actually impact Application Performance ?
5. Internet Traffic Engineering for performance relies on probing ? What is the reliability of probing in the Internet ? How to improve it ?
6. What about those customers who can’t get or won’t pay for BGP Multi-homing ? DNS, NAT, Shim6, ...
7. What about uncontrolled frequent BGP updates for ingress Inter-domain Traffic Engineering ? Is route dampening enabled ? If not, why not ? If so, why the concern ?

