

Impact of BGP Dynamics on Intra-Domain Traffic and Router Load

Sharad Agarwal

w/ C.N. Chuah, S. Bhattacharyya, C. Diot




Internet Architecture

- Autonomous Systems
 - Separately administered {routers, links, hosts}
 - 15,000
 - Tier-1 ISP (global / national backbone)
 - Regional ISP
 - Small stub AS (company, university)
- Interaction between ASes
 - Internal topology not shared
 - Exchange reachability

2

Border Gateway Protocol

- Path vector protocol

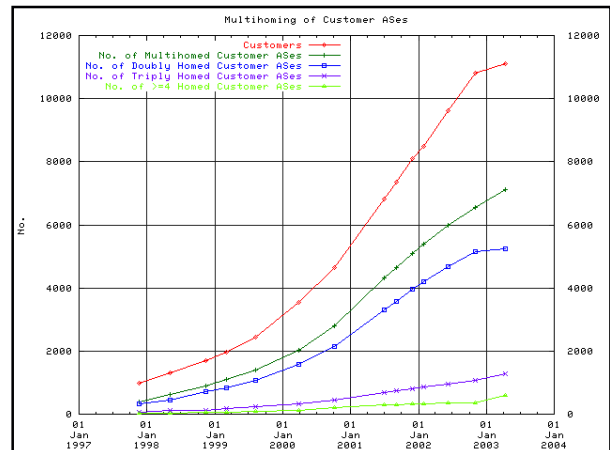
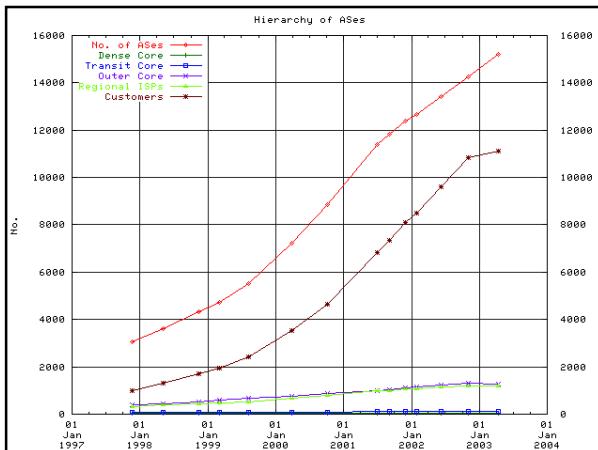


3


BGP

- Policy routing
 - Receive many announcements
 - Import policy
 - Path selection policy
 - Export policy
 - Configured by operator
- ~15 years old
 - Performing to original design expectations?

4



BGP Routing Table Changes



7

The Sky is Falling?

- BGP concerns
 - [Huston01], [Maenell02], [Cowie01] ...
 - Too many updates
 - Too many routing table entries
 - Too much multihoming
- Is this bad?

8

Overview

- Problems?
 - Router load
 - Intra-ISP traffic shifts
 - Failover
 - Traffic balancing
 - Bogus announcements
 - Static peering
 - OPCA solves these

9

Data

- BGP data
 - Route updates
 - Sprint (RRCs, eBGPs, RR mesh)
- SNMP data
 - Router CPU utilization
 - Sprint
- Packet traces
 - 44 bytes via optical splitters
 - Sprint (ingress, egress, backbone)


10

Problem #1: Router Load

- Many BGP updates
 - Too much traffic? It's peanuts!
 - 100 update/sec, 4096B max => 400KB/s
 - Router processing?
 - Every update
 - In-filtering, route selection, FIB updates, out-filtering, RIB updates
 - Overloads notoriously cause router shutdown
 - Cisco routers (200 Mhz R5000, 256-512MB)


11

BGP Activity & CPU Load 10-17 Feb 2003




12

SQL Worm – 21:30 Fri 24 Jan 2003



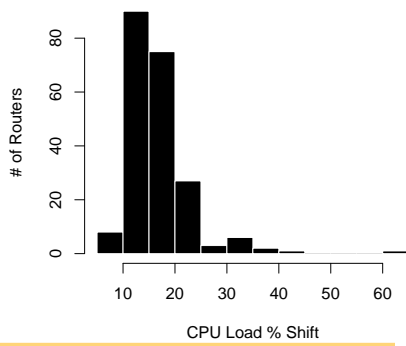
13

Histogram of CCF over 1 hour



14

Histogram of CPU Load Shift over 3 days



15

Summary

- High router CPU utilization
 - Can increase convergence times
 - Can cause outage
- Some weak correlation exists on average
- Specific incidents have high correlation
 - But shift in CPU load is small
- **BGP activity hasn't impacted router loads**
 - At long time scales O(5 mins)

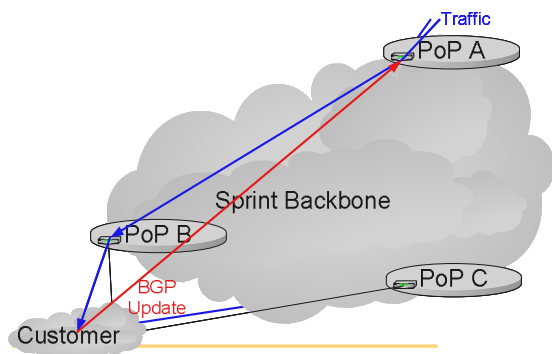
16

Overview

- Problems?
 - Router load
 - Intra-ISP traffic shifts

17

Example: Shift in Intra-Domain Traffic



18

Problem #2: Traffic Shifts

Do BGP updates change how traffic traverses Sprint's network?

- Why is this important?
 - Latency variation affects applications, e.g. VoIP
 - Unstable traffic matrix makes provisioning / traffic engineering harder
 - e.g. IGP weights

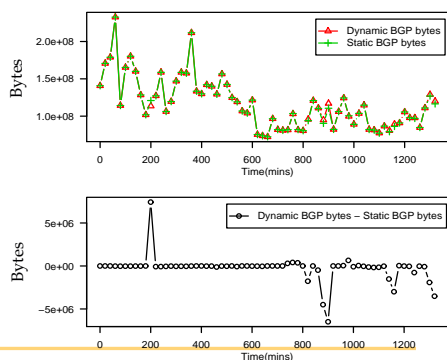
19

Traffic Analysis Methodology

- 22 hour pkt trace, >2 billion pkts & iBGP data
 - Find egress PoP for each data packet
 - Longest prefix match, router to PoP map
 - Ingress link to multiple egress PoP fan-out
 - Identify traffic variability due to BGP updates
 - Static BGP table + data packets
 - Shows variability due to other factors
 - Dynamic BGP table + data packets

20

Variability in Traffic to an Egress PoP



Summary of Results

- <0.03% of ingress traffic shifts egress PoP
 - Representative of all 5 traces
- Specific ASes are sources of ~50% of variability
- Shifts happen <3 times for any flow
- Why?
 - Elephants prevalent
 - ~1.5 % of 200,000 dest addrs → 80% traffic
 - 0-15 updates / minute
 - NH/Prefix = 0.11 versus 0.63 for all

22

Summary

- BGP updates hardly affect intra-Sprint traffic fan-out
 - AT&T[Rexford02]: stable traffic → stable prefixes
 - Why?
 - Standard route filtering?
 - Stable prefixes attract stable traffic?
 - Less latency variation for applications
 - Traffic matrix more stable
 - Good for TE, provisioning

23

Overview

- Problems?
 - Router load
 - Intra-ISP traffic shifts
 - Failover
 - Traffic balancing
 - Bogus announcements
 - Static peering
 - OPCA solves these

24