



BGP ANYCAST Simulations Using GTNetS (work in progress)

Talal Jaafar

Georgia Tech & CAIDA



Agenda

- IP Anycasting Studies
- Our Approach
- GTNetS Simulations
- Experiments
- Remaining Work
- Q&A



IP Anycast Studies

- Current measurements
 - Planet Lab measurements (Sarat, Terzis et al)
 - Decrease in latency with anycast
 - Clients don't always hit nearest server
 - Relatively small number of outages, but lasted for long time (>100s)
 - Hint more global nodes might cause instability
 - Ripe K-root measurements (Lorenzo)
 - Good latency
 - Local nodes take load off global nodes, but not by a great factor
 - Quite stable – few switches
 - Ballani and Francis with their Planet Lab measurements conclude that anycast nodes are quite stable – hardly any flips observed



Why Simulations?

- Study the impact of BGP on Anycast
 - BGP Convergence
 - Flap Dampening
- Study the impact of Anycasting on BGP
 - BGP Churn
 - Large # of global nodes → Convergence impact?
- Simulations might be useful in analyzing different options for good placements of future anycast servers



GTNetS Simulation

- Discrete-Event Packet Level simulations
- BGP : BGP++ implementation ported from ns-2 (zebra based)
- Simulation handles actual routing – longest prefix match based FIB which is populated by BGP
- Anycast servers supported using /32 prefix address advertisement



GTNetS Simulations (cont'd)

- Failures
 - Silent link failure
 - Explicit withdraw of Anycast prefix
- Modes of Deployment
 - Global nodes
 - Hierarchical Local-Global nodes
 - Hierarchical Local-Global nodes with Local nodes advertising more specific prefix
- Topology
 - Real AS Topology (RouteViews Project)
 - 2 Stages



Tier-1 Topology Simulation

- Representative f-root Internet backbone connections inferred from routeviews.
 - Simulate 1 BGP speaker per AS
 - Peer-Peer, Customer and Provider relations inferred and appropriate policies applied.
 - Total of 44 large ISPs with 467 interconnecting links simulated



Tier-1 Topology Simulation (3)

- Simulated Topology:
 - 10 Anycast servers (All Global)
 - 9 in North America
 - 1 in Europe
 - 34 clients
 - 22 in North America
 - 12 in Europe and Asia

Tier-1 Topology Simulation Results

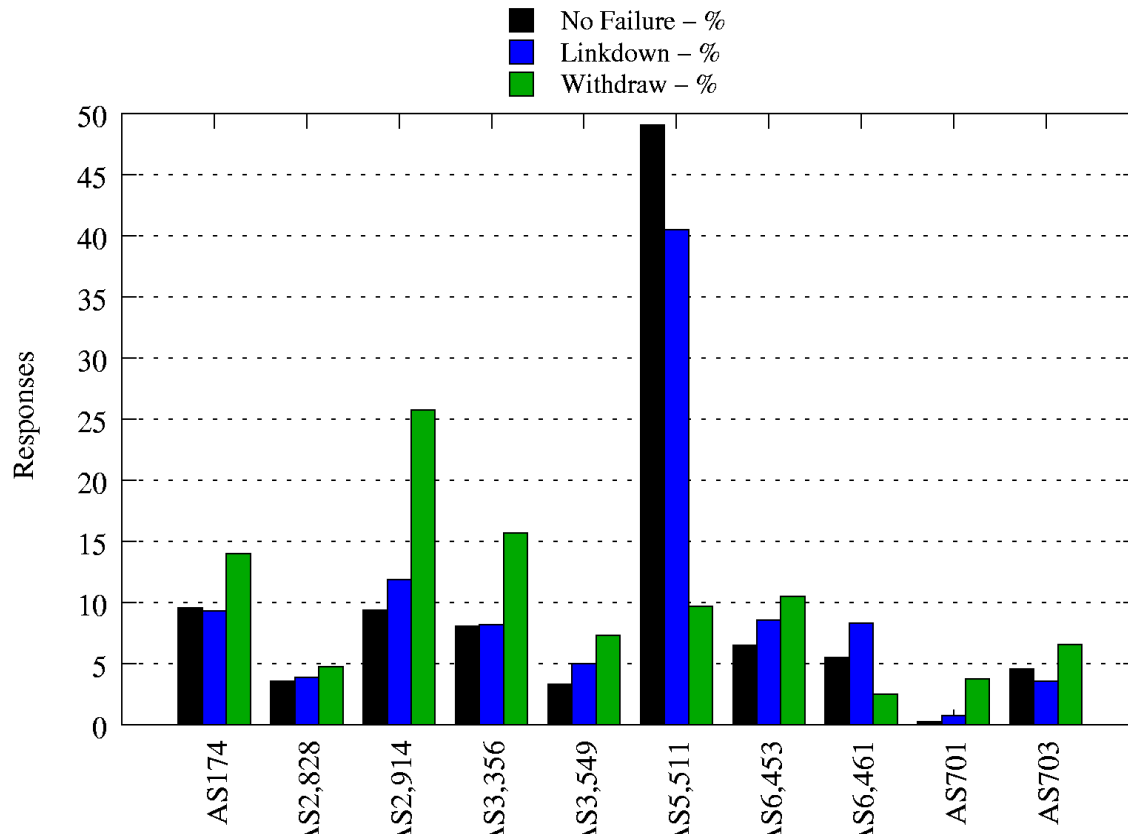


Figure 1: Load distribution (44 Node Topology)

Tier-1 Topology Simulation Results




Figure 1: Geographic load distribution (44 Node Topology)



Tier-1 and Tier-2 Topology Simulation

- 5,476 Nodes and 14,468 Links
- 12 Nodes provides service to F-Root
- Distributed Simulation



Research Goals

- Use Simulations to:
 - Compare Unicast Vs. Anycast Routing
 - No failures
 - Link failures
 - Prefix withdrawals
- BGP Convergence
BGP Churn
Anycast Server Stability
Anycast Load Balancing



Questions

?