Eliminating Packet Loss Caused by BGP Convergence

Nate Kushman

Srikanth Kandula, Dina Katabi, and Bruce Maggs

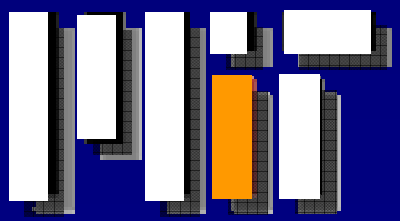

The Problem:
BGP Convergence Causes You Packet Loss

- Route changes cause up to 30% packet loss for more than 2 minutes *[Labovitz00]*
- Even for domains dual homed to tier 1 providers, a failover event can cause multiple loss bursts, and one loss burst can last for up to 20s *[Wang06]*
- Popular and unpopular prefixes experience losses due to BGP convergence *[Wang05]*
- 50% of VoIP disruptions are highly correlated with BGP updates *[Kushman06]*

What Kind of Solution Do We Want?

- Eliminate packet loss during BGP convergence
- An adopting ISP protects itself and its customers from loss even if no other ISP cooperates

Transient Path Loss Problem

All of Patrick's providers are using him to get to MIT

Nobody offers Patrick an alternate path

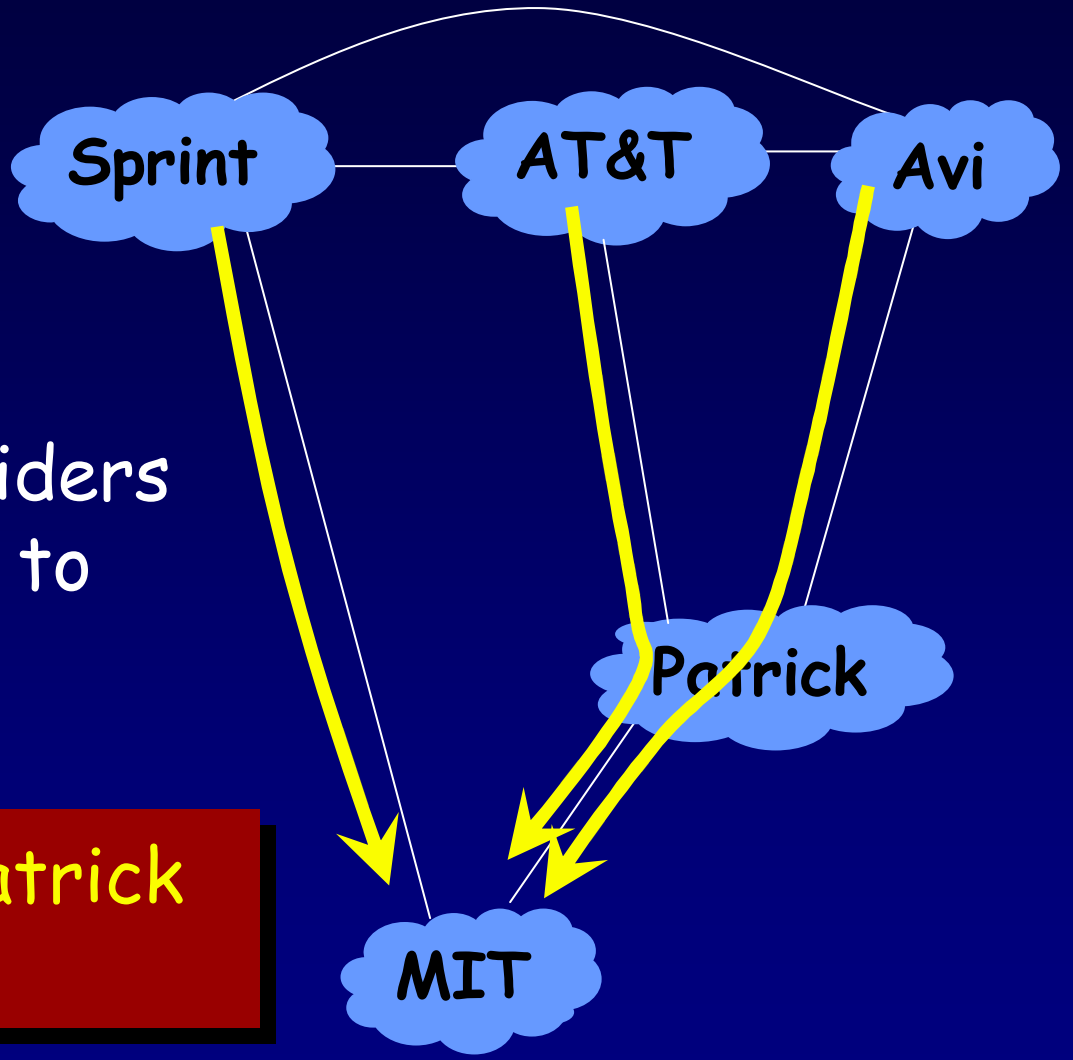

Transient Path Loss Problem

Patrick knows no alternate path to MIT

Patrick drops AT&T's and Avi's packets to MIT, and his own

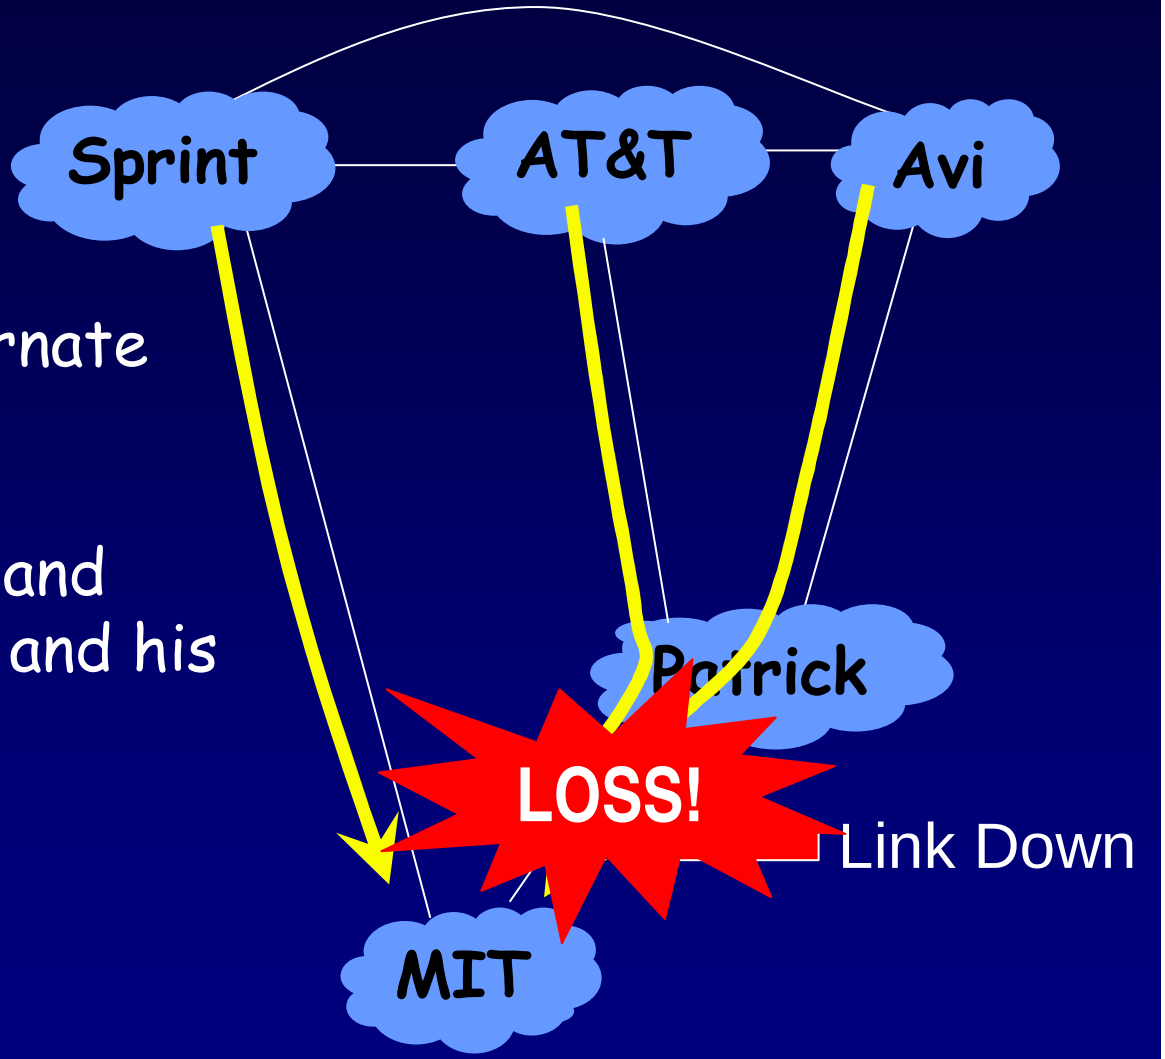

Transient Path Loss Problem

Eventually, Patrick withdraws path from AT&T and Avi

AT&T and Avi stop sending packets to Patrick

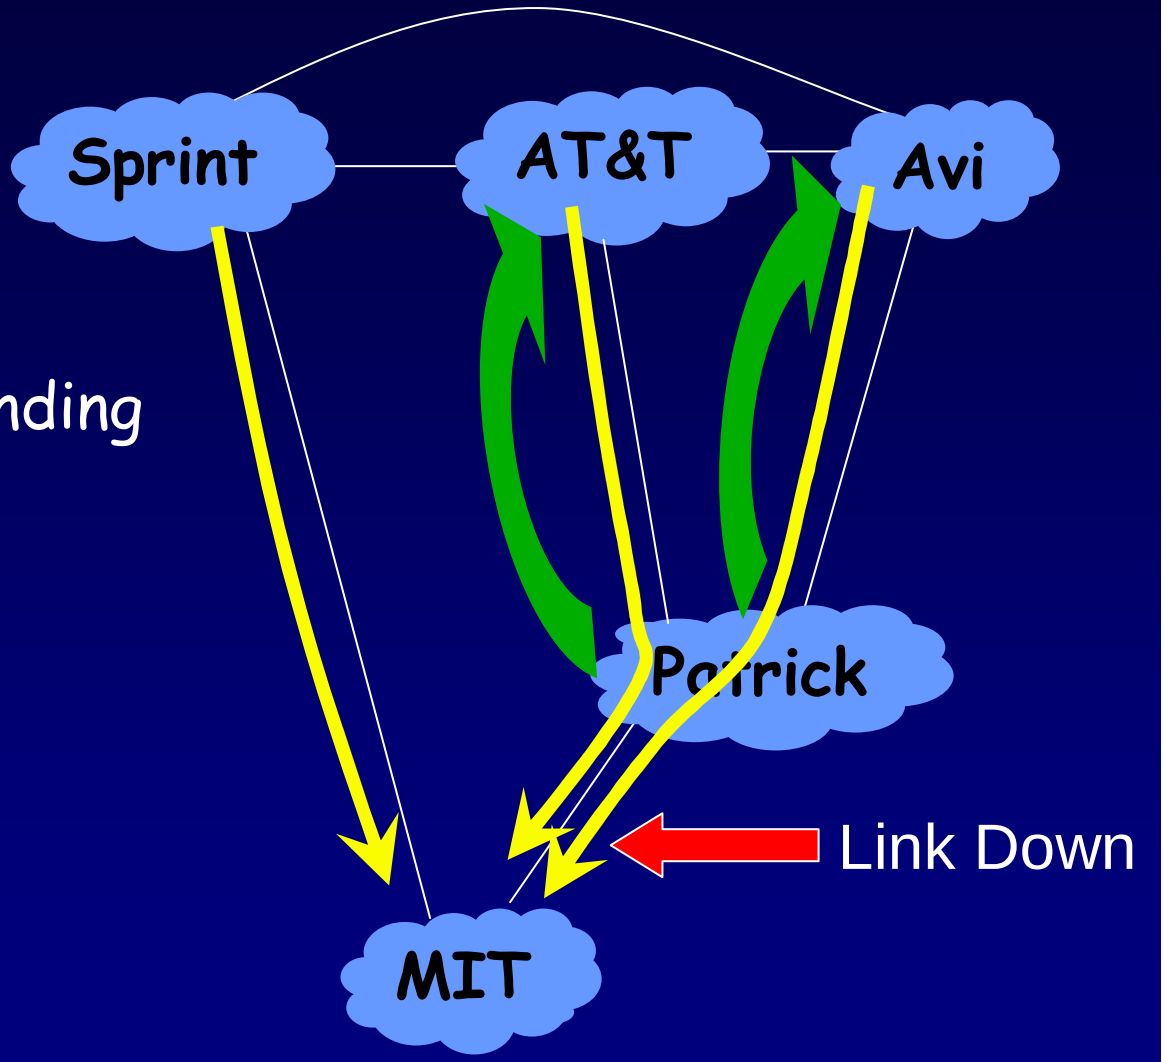

Transient Path Loss Problem

Eventually, Patrick withdraws path from AT&T and Avi

AT&T and Avi stop sending packets to Patrick

AT&T announces the Sprint path to Patrick → Traffic flows

Temporary Packet Loss

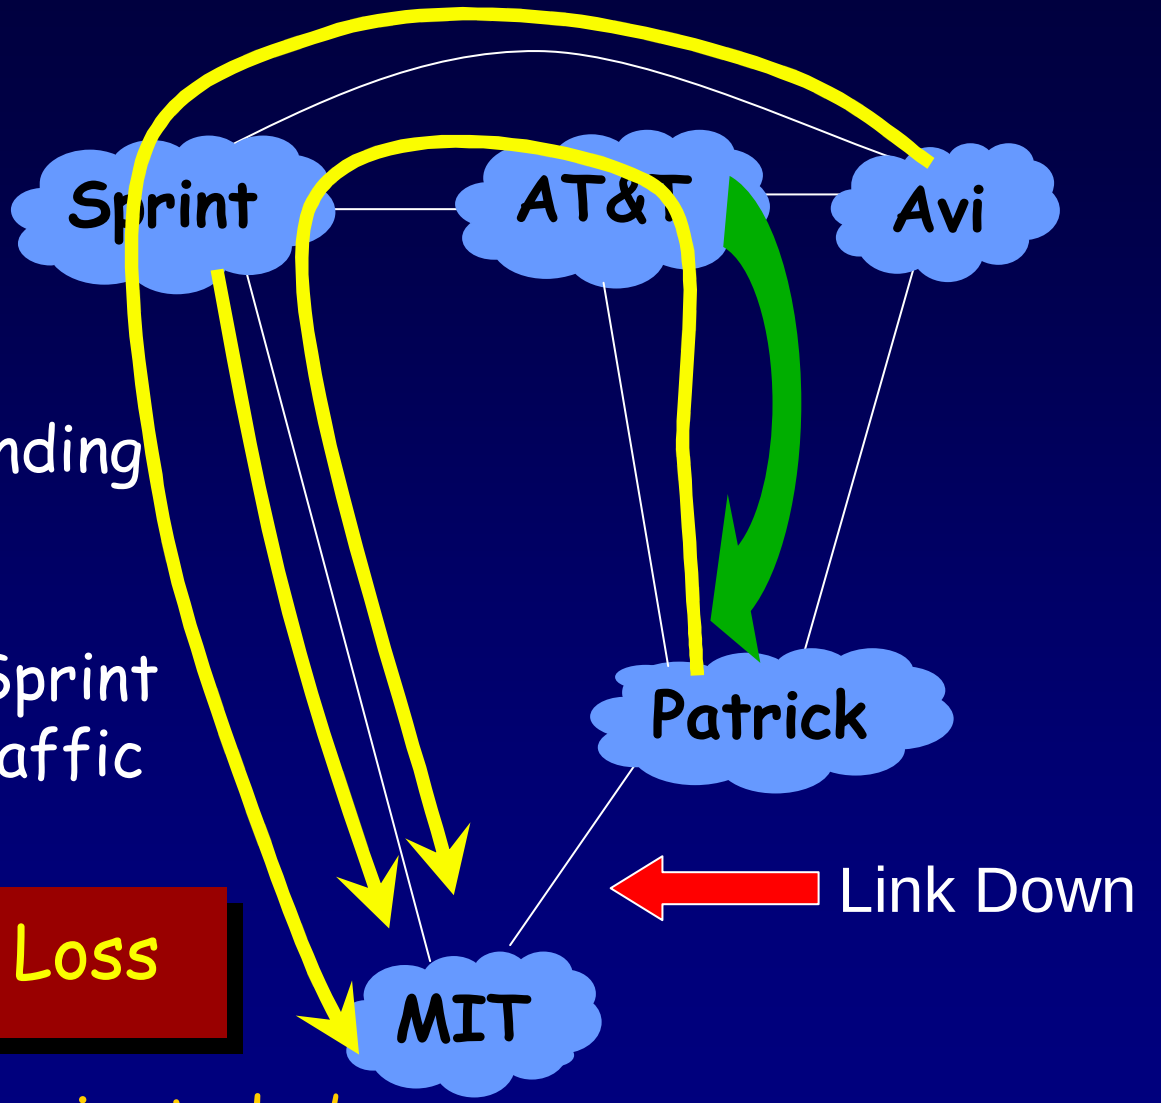

Significant loss happens in today's Internet, even when connected to Tier 1s

How do we solve Patrick's problem?

Tell Patrick a failover path before the link fails

rather than after it, as is often the case in current BGP

Help Patrick Help You!

AT&T advertises to Patrick "AT&T→ Sprint → MIT" as a failover path

Link Fails → Patrick immediately sends traffic on failover path

No Loss !

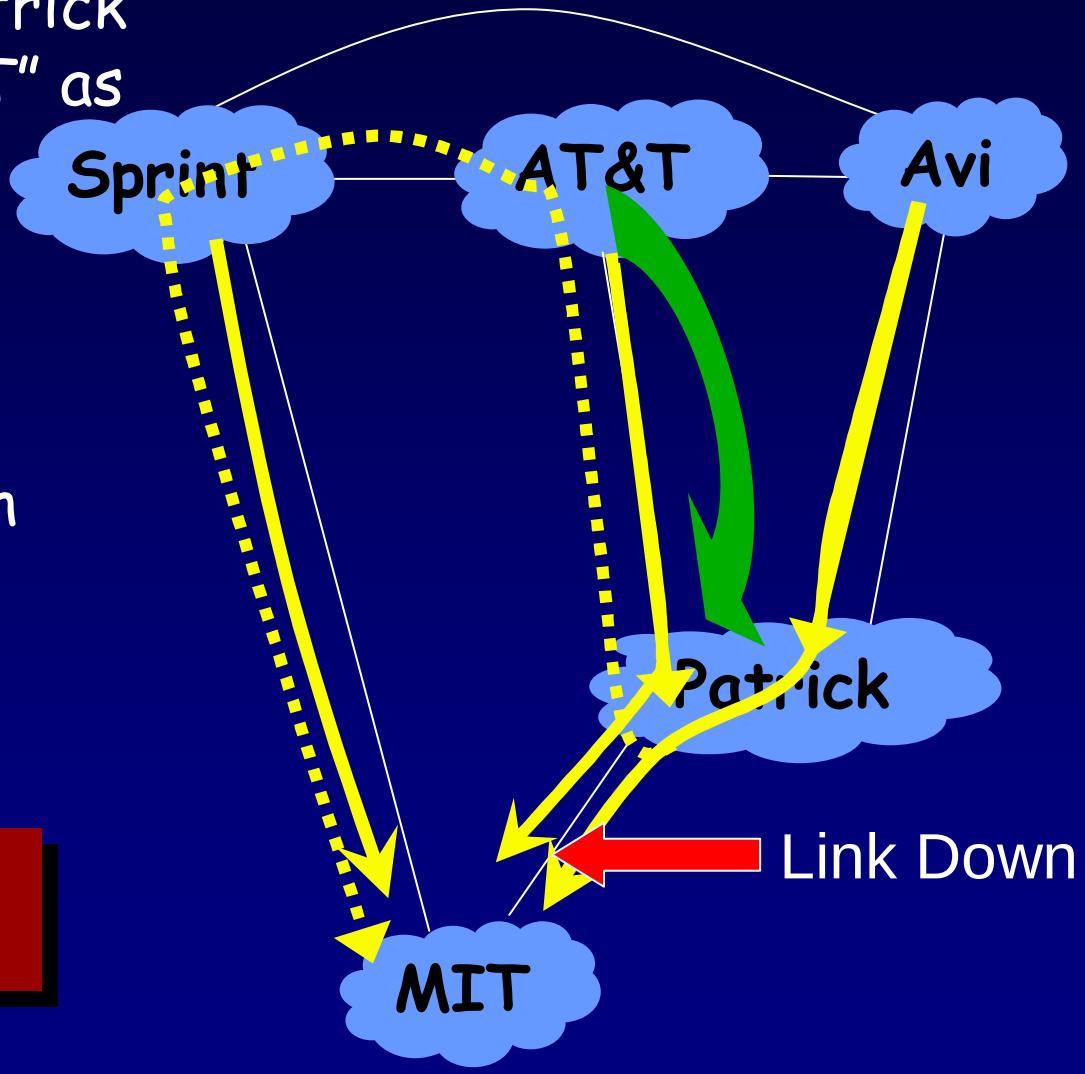

Our Solution: Two Simple Rules

Routing Rule: Each router advertises only one failover path and only to the next hop router on its primary path

Forwarding Rule: When routers receive packets from the next-hop interface for their primary path they forward them along the failover path

Guarantee: A router is guaranteed to see no BGP-caused packet loss during convergence, if it will have a valley-free path to the destination at convergence

Helps Even When Deployed in a Single AS

Currently Patrick drops packets even if he knows an alternate path

R1 offers Failover path to R2

R2 U-turns packets – back to R1

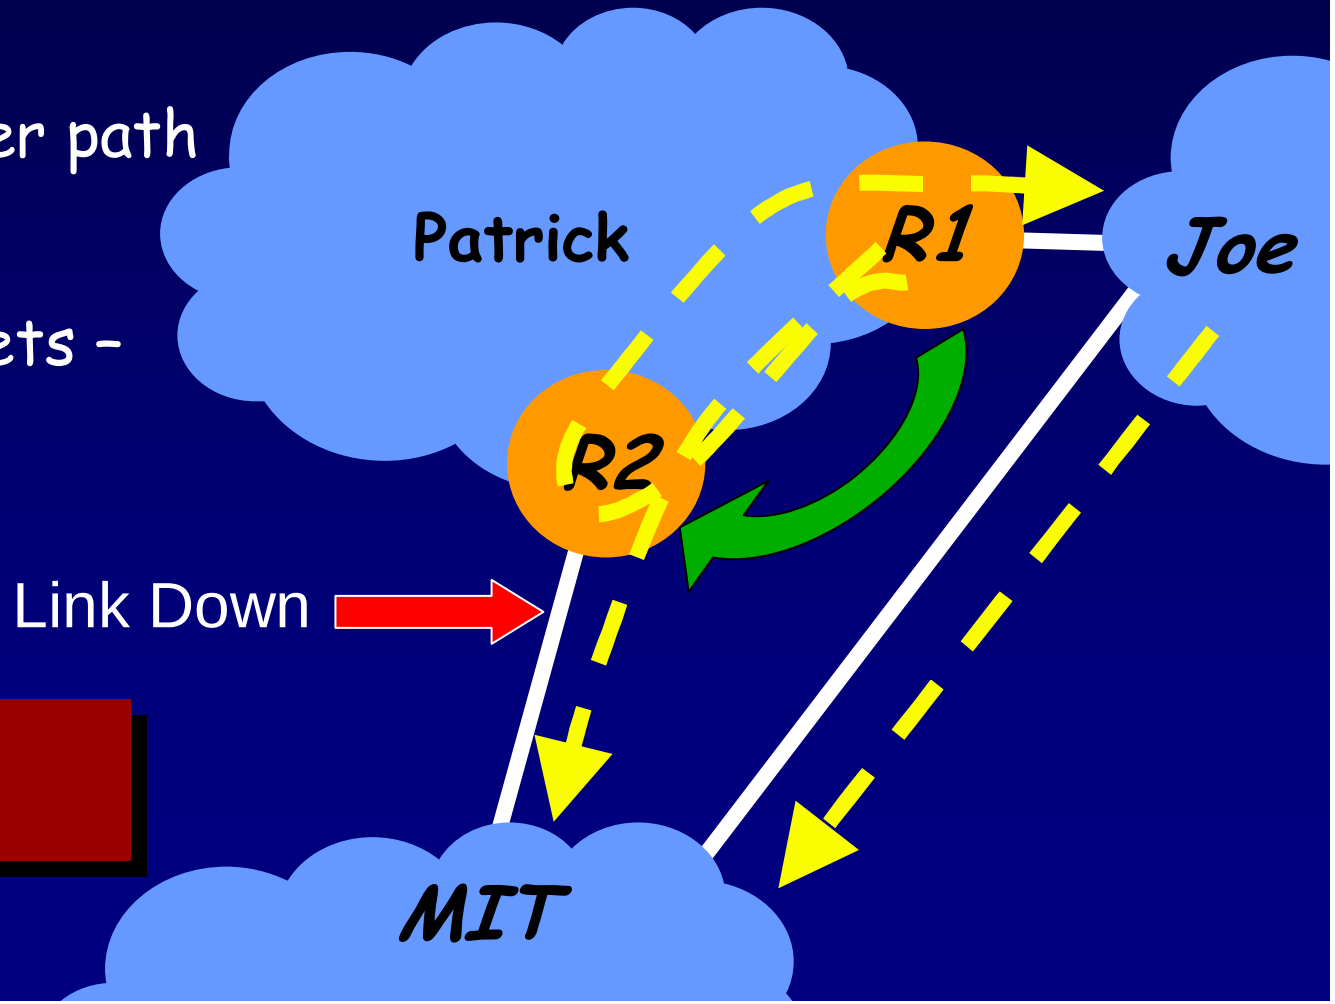

No Loss !

Experimental Results

- Router-Level Simulation over the full Internet
  - AS-graph from Routeviews and RIPE BGP Data
  - Use inference algorithms to annotate links with customer-provider or peer relationships
  - Add border routers based on the connections to other AS
  - Used internal MRAI of 5s and external MRAI of 30s
- For each experiment:
  - Random destination
  - Take down a Random Link
  - Find the duration of packet loss for routers using the down link which have a path after convergence
  - Run for 1000 Randomly Chosen Links and Destinations

Significant Benefit Running Only in AT&T

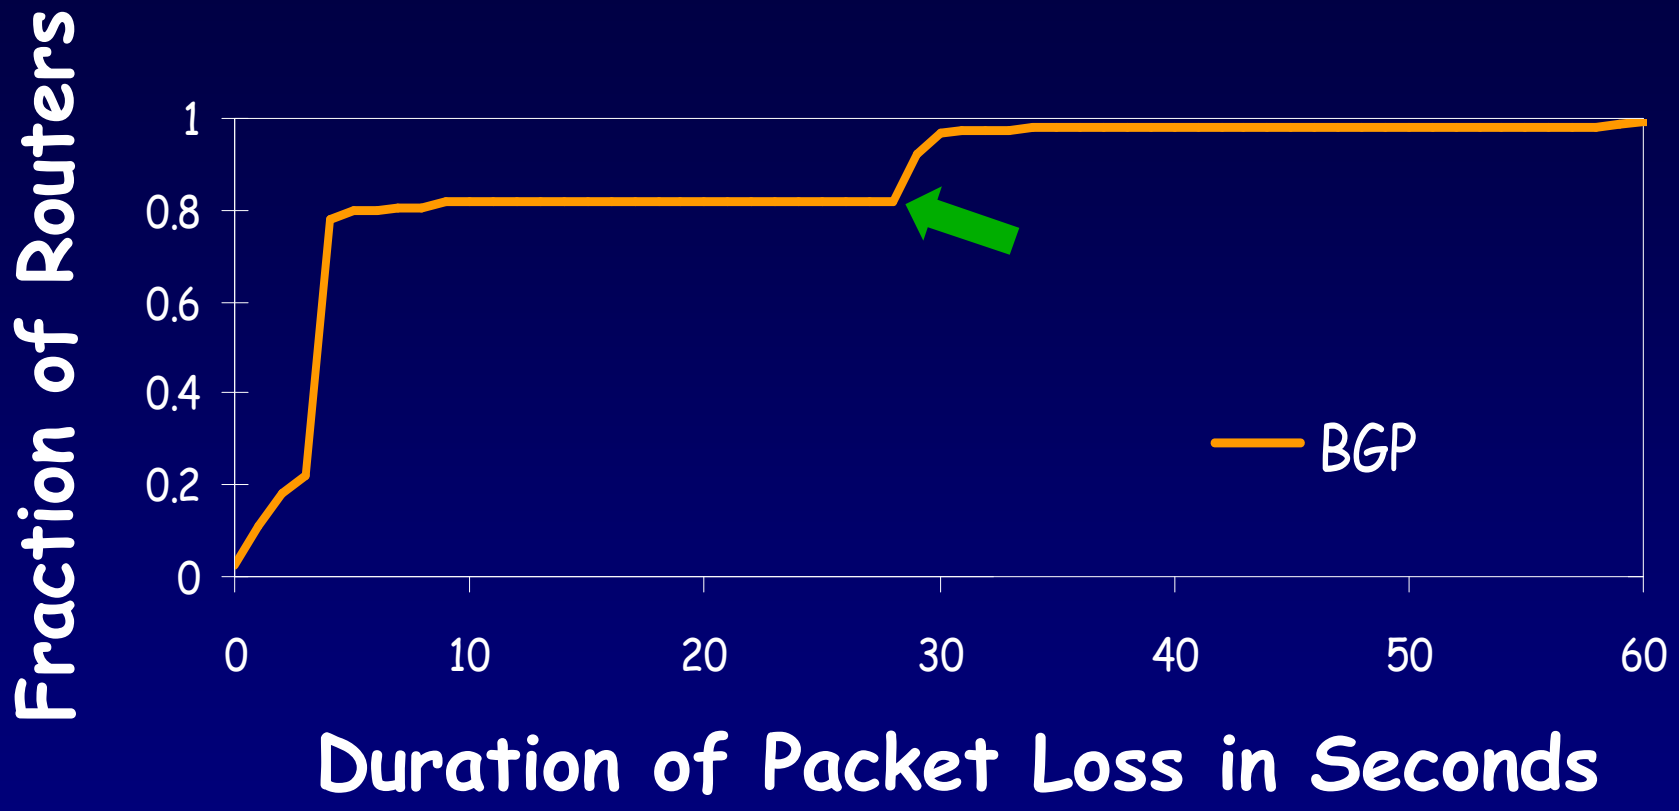

~20% See 30s or More of Packet Loss

Significant Benefit Running Only in AT&T

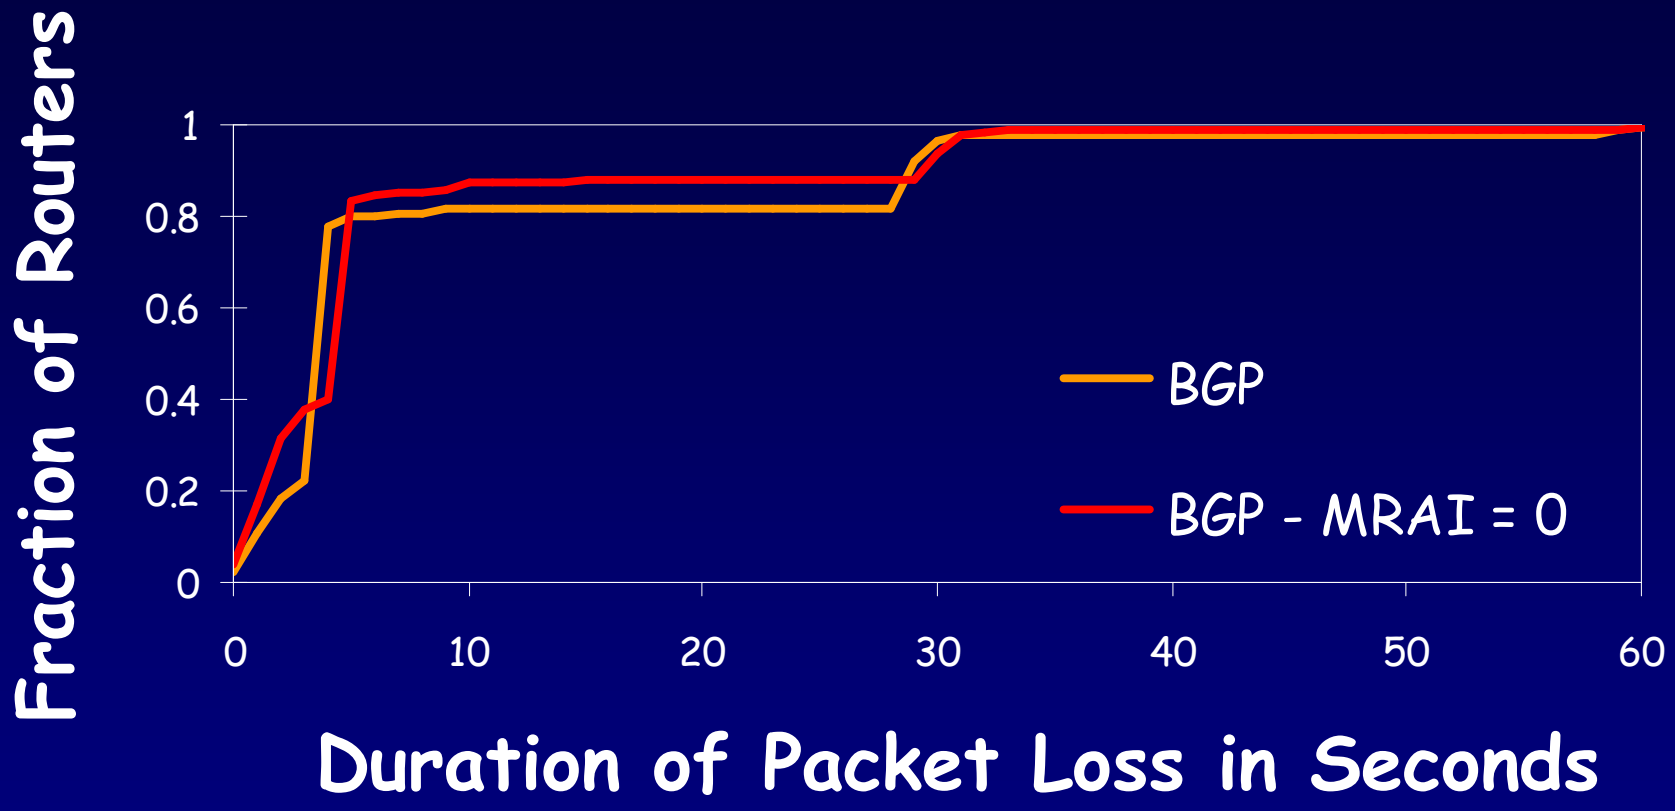

Setting MRAI to 0 still leaves Significant Packet Loss

Twice the Number of Updates for Both AT&T and Customers

Significant Benefit Running Only in AT&T

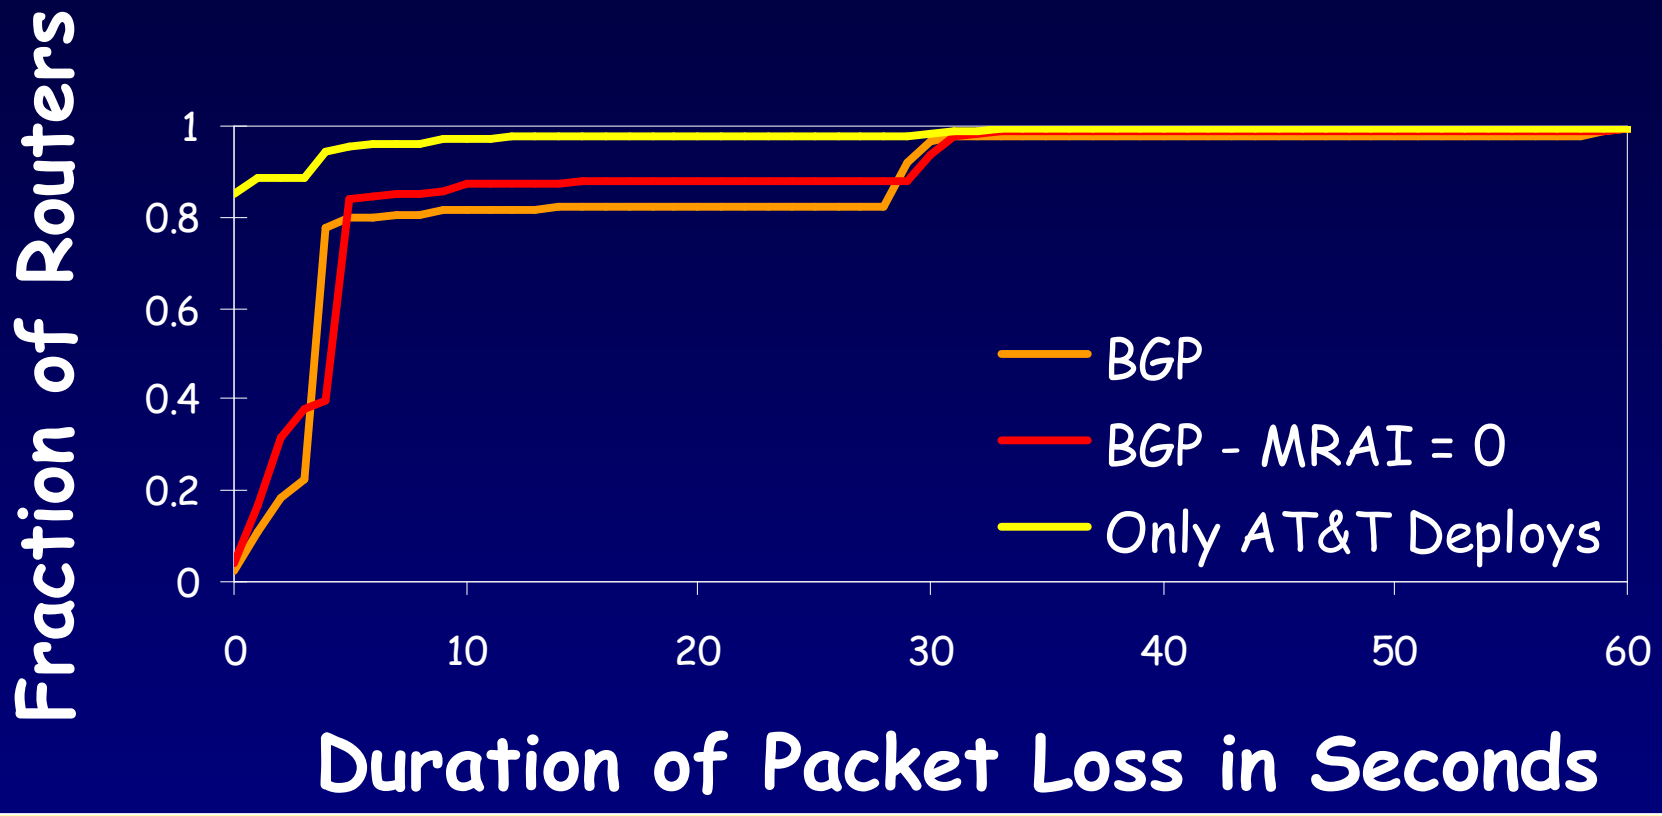

Less than 3% if only AT&T adopts

Full Benefit Once Running Everywhere

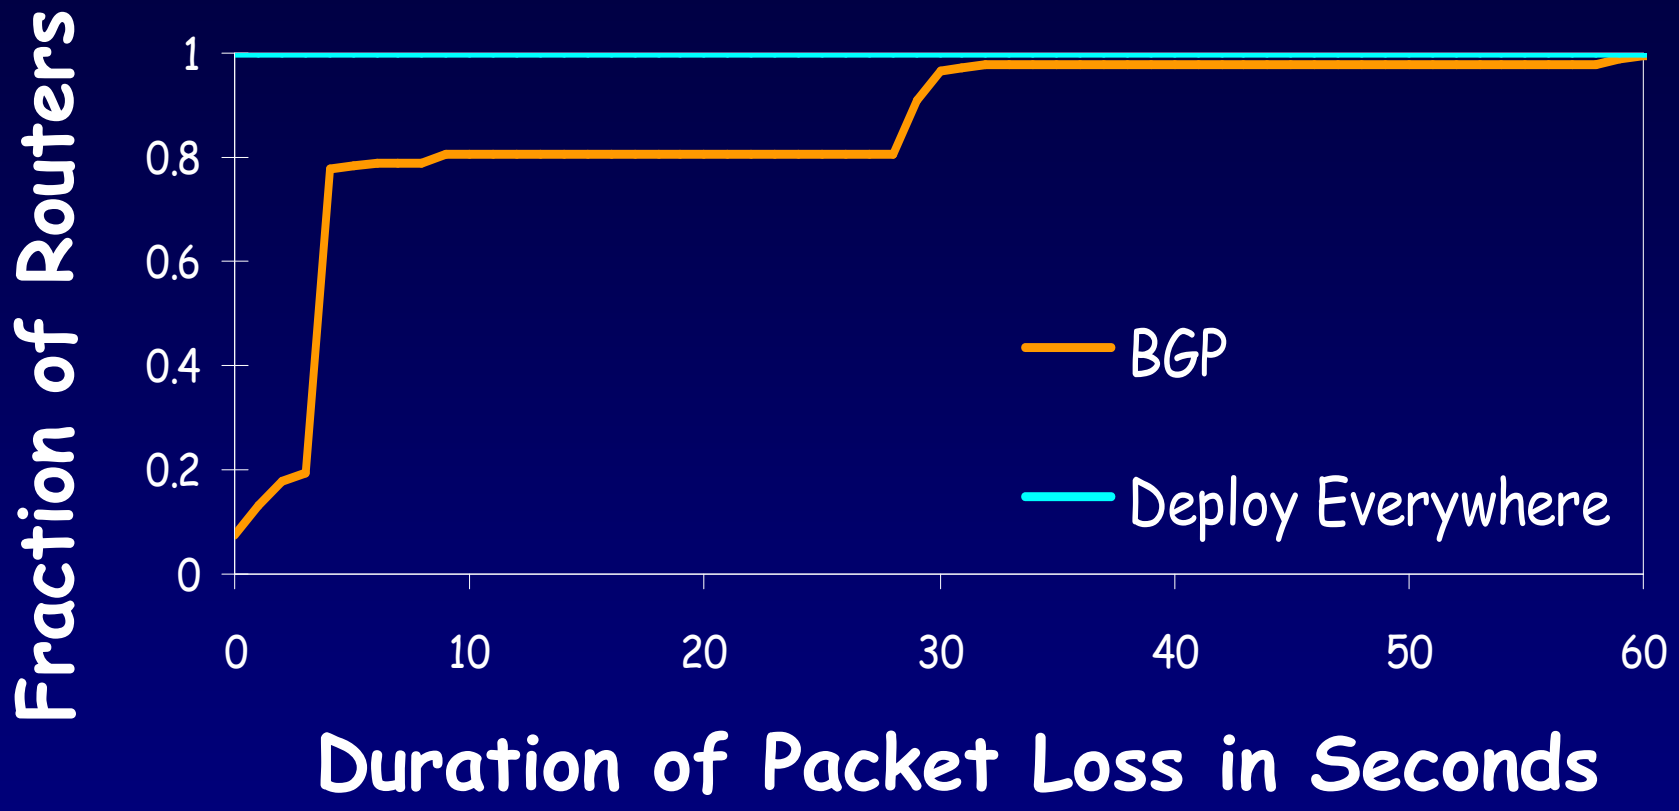

Running Everywhere Eliminates All Packet Loss

Little Additional Overhead

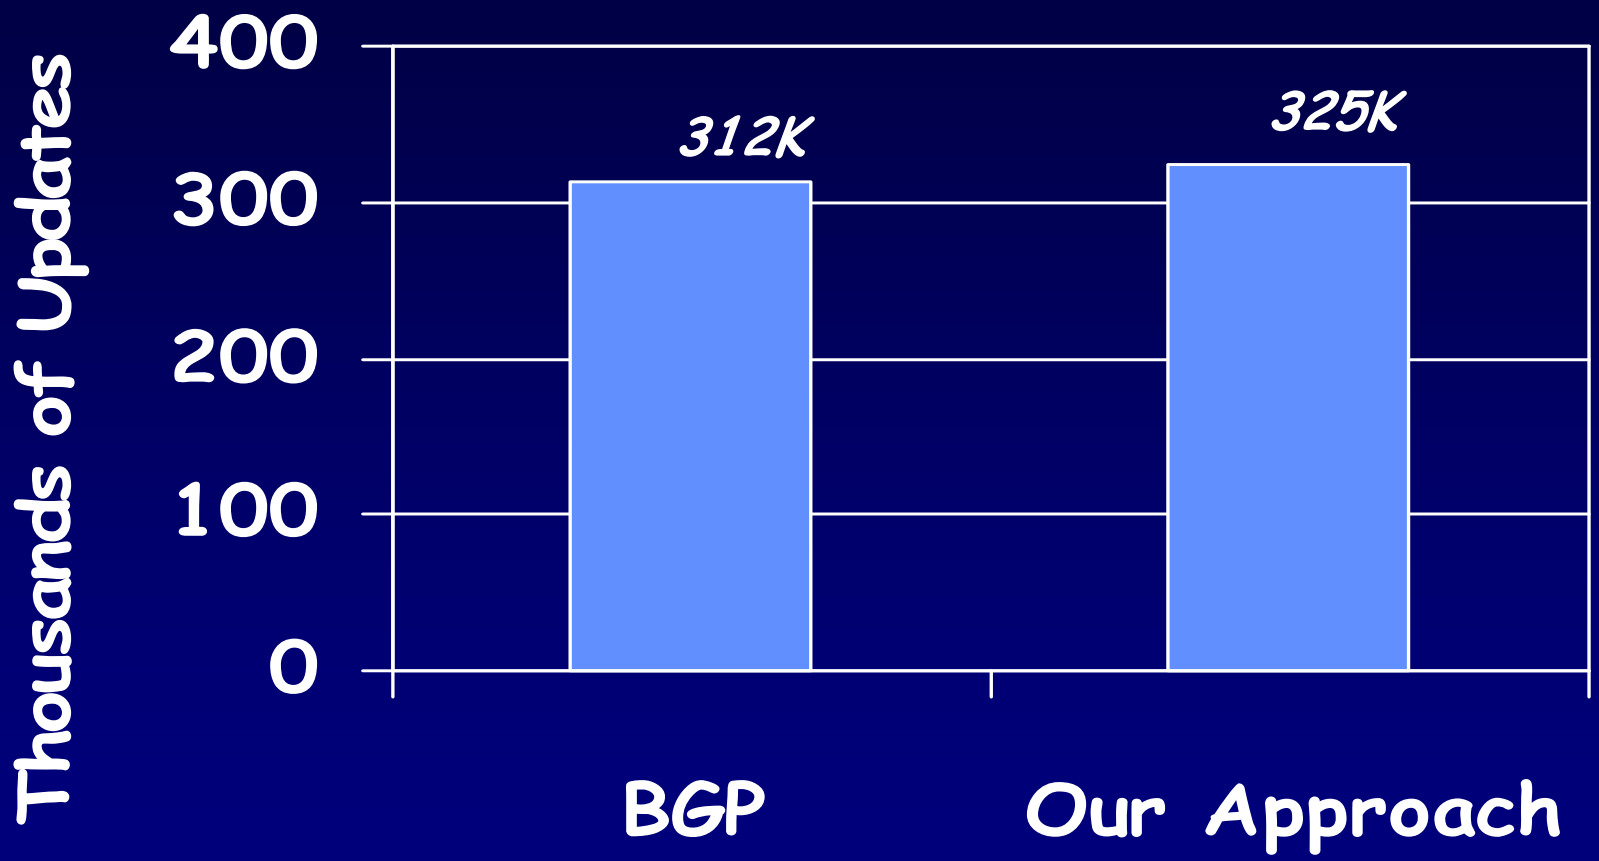

Less than 5% more updates network wide

Conclusion

- Simple Mechanism
- Solves Problem
- Deployable

- Offer a failover path only to next-hop neighbor
- Eliminates packet loss resulting from BGP convergence
- An adopting ISP reduces loss even if no other ISP cooperates