SIGCOMM 2011
Toronto, ON
Aug. 16, 2011

Let the Market Drive Deployment

A Strategy for Transitioning to BGP Security

Phillipa Gill
University of Toronto

Michael Schapira
Princeton University

Sharon Goldberg
Boston University

Incentives for BGP Security

Insecurity of Internet routing is well known:

- **S-BGP** proposed in 1997 to address many issues
- Challenges are being surmounted:
  - Political: Rollout of RPKI as a cryptographic root trust
  - Technical: Lots of activity in the IETF SIDR working group

The pessimistic view:

- This is economically infeasible!
- Why should ISPs bother deploying **S*BGP**?
- No security benefits until many other ASes deploy!
- Worse yet, they can't make money from it!

Our view:

- Calm down. Things aren't so bad.
- ISPs **can** use S*BGP to make money
- …by attracting traffic to their network.

Outline

- Part 1: Background
- Part 2: Our strategy
- Part 3: Evaluating our strategy
  - Model
  - Simulations
- Part 4: Summary and recommendations

Traffic Attraction & Interception Attacks

April 2010 : China Telecom intercepts traffic

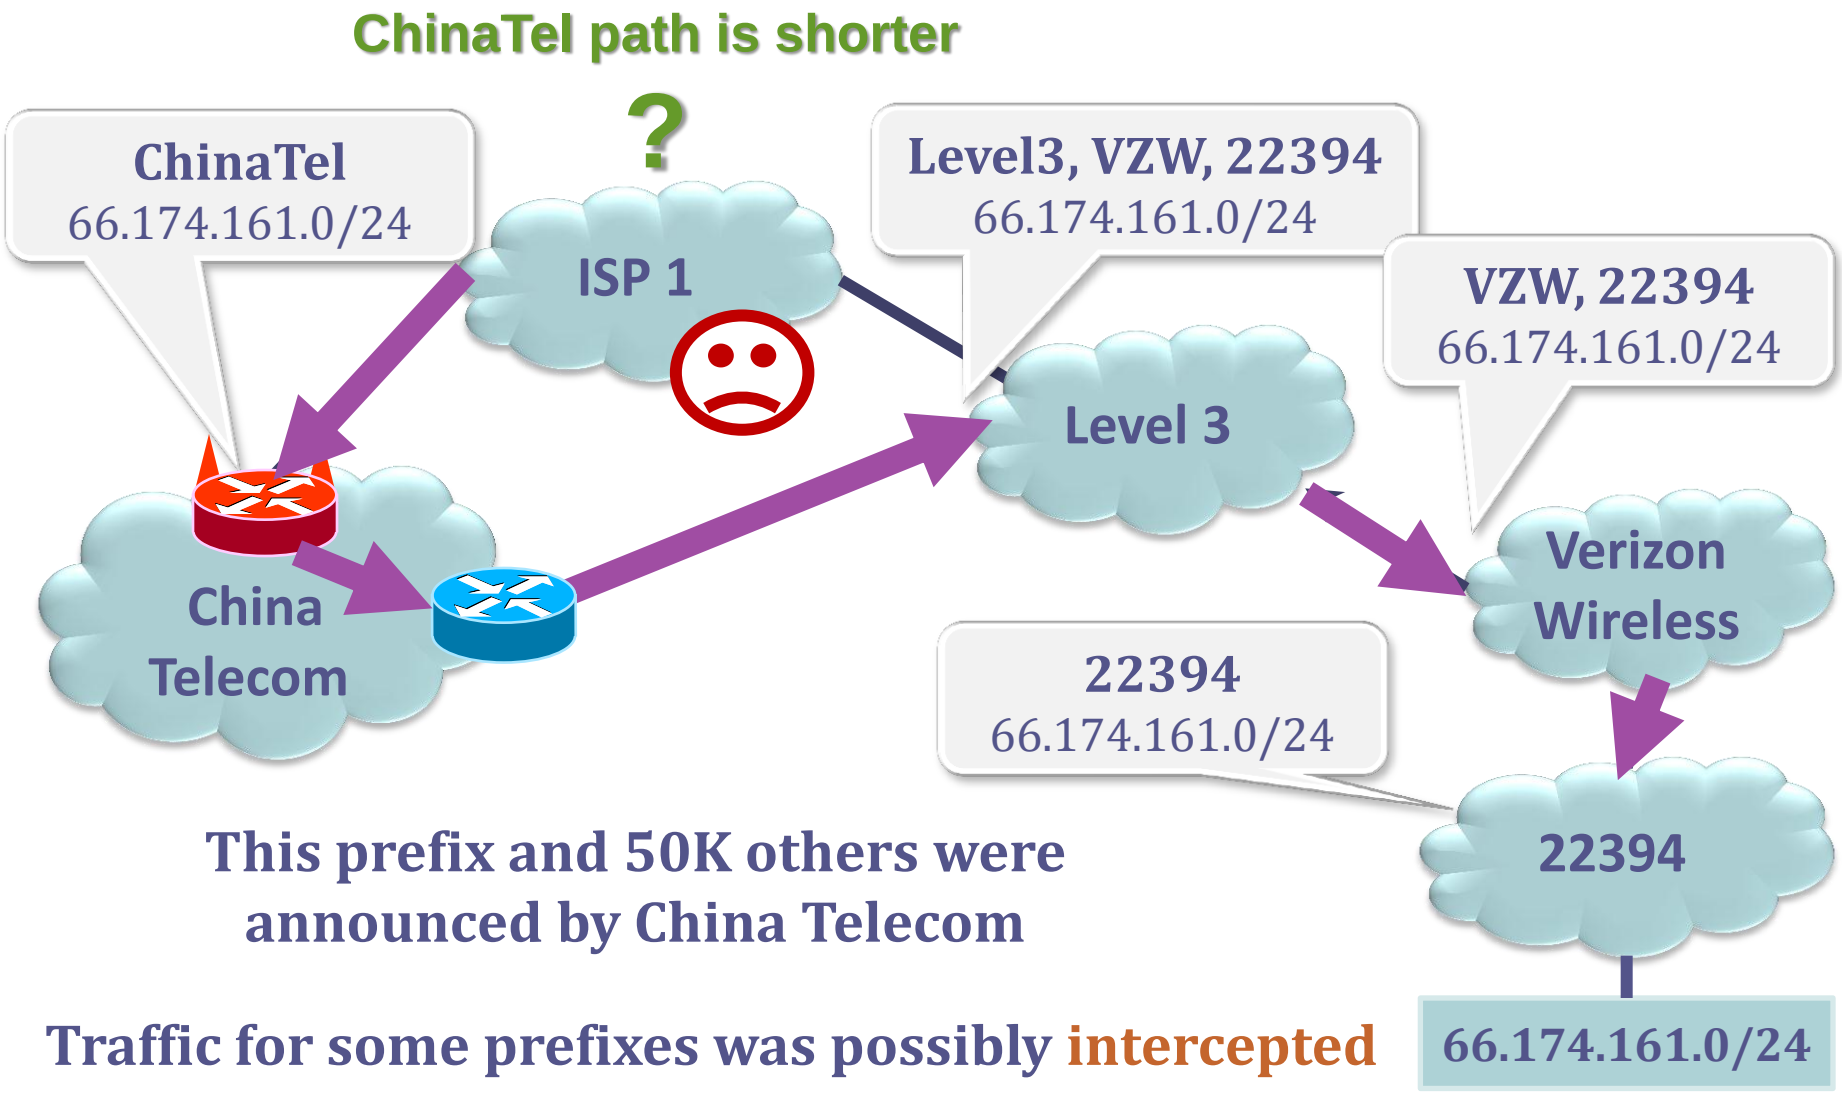

This prefix and 50K others were announced by China Telecom

Traffic for some prefixes was possibly intercepted

Securing the Internet: RPKI

Resource Public Key Infrastructure (RPKI): Certified mapping from ASes to public keys and IP prefixes.

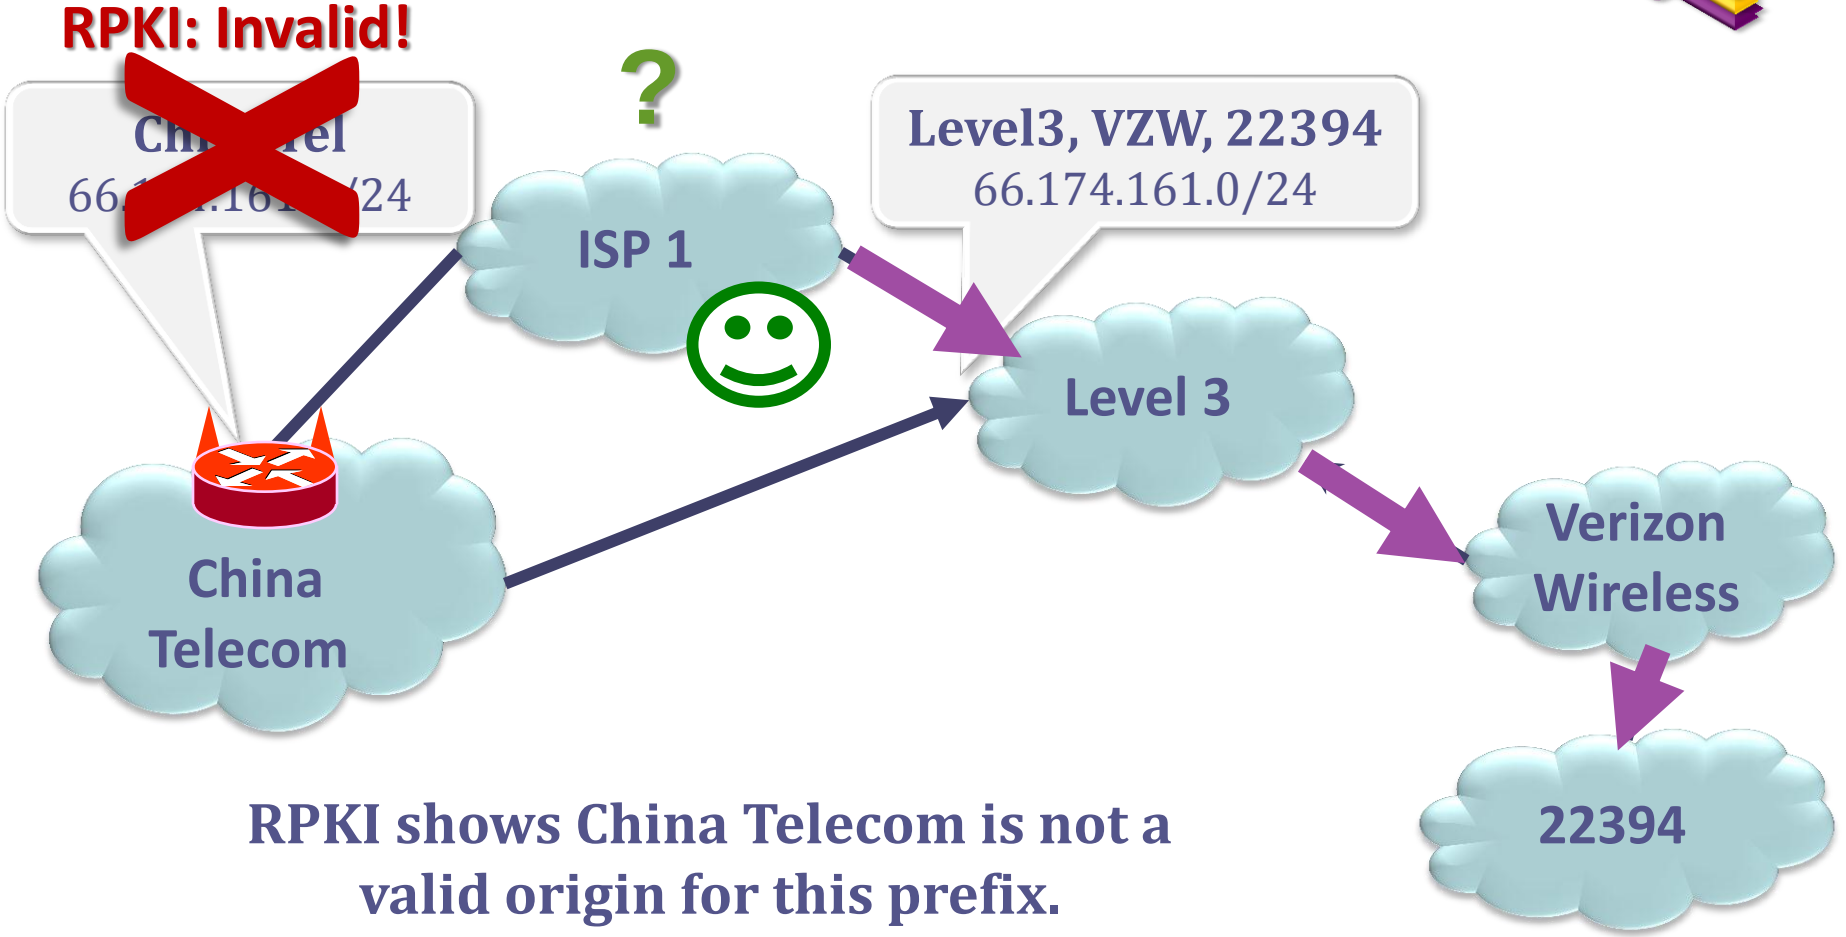

RPKI shows China Telecom is not a valid origin for this prefix.

But RPKI alone is not enough!

Resource Public Key Infrastructure (RPKI): Certified mapping from ASes to public keys and IP prefixes.

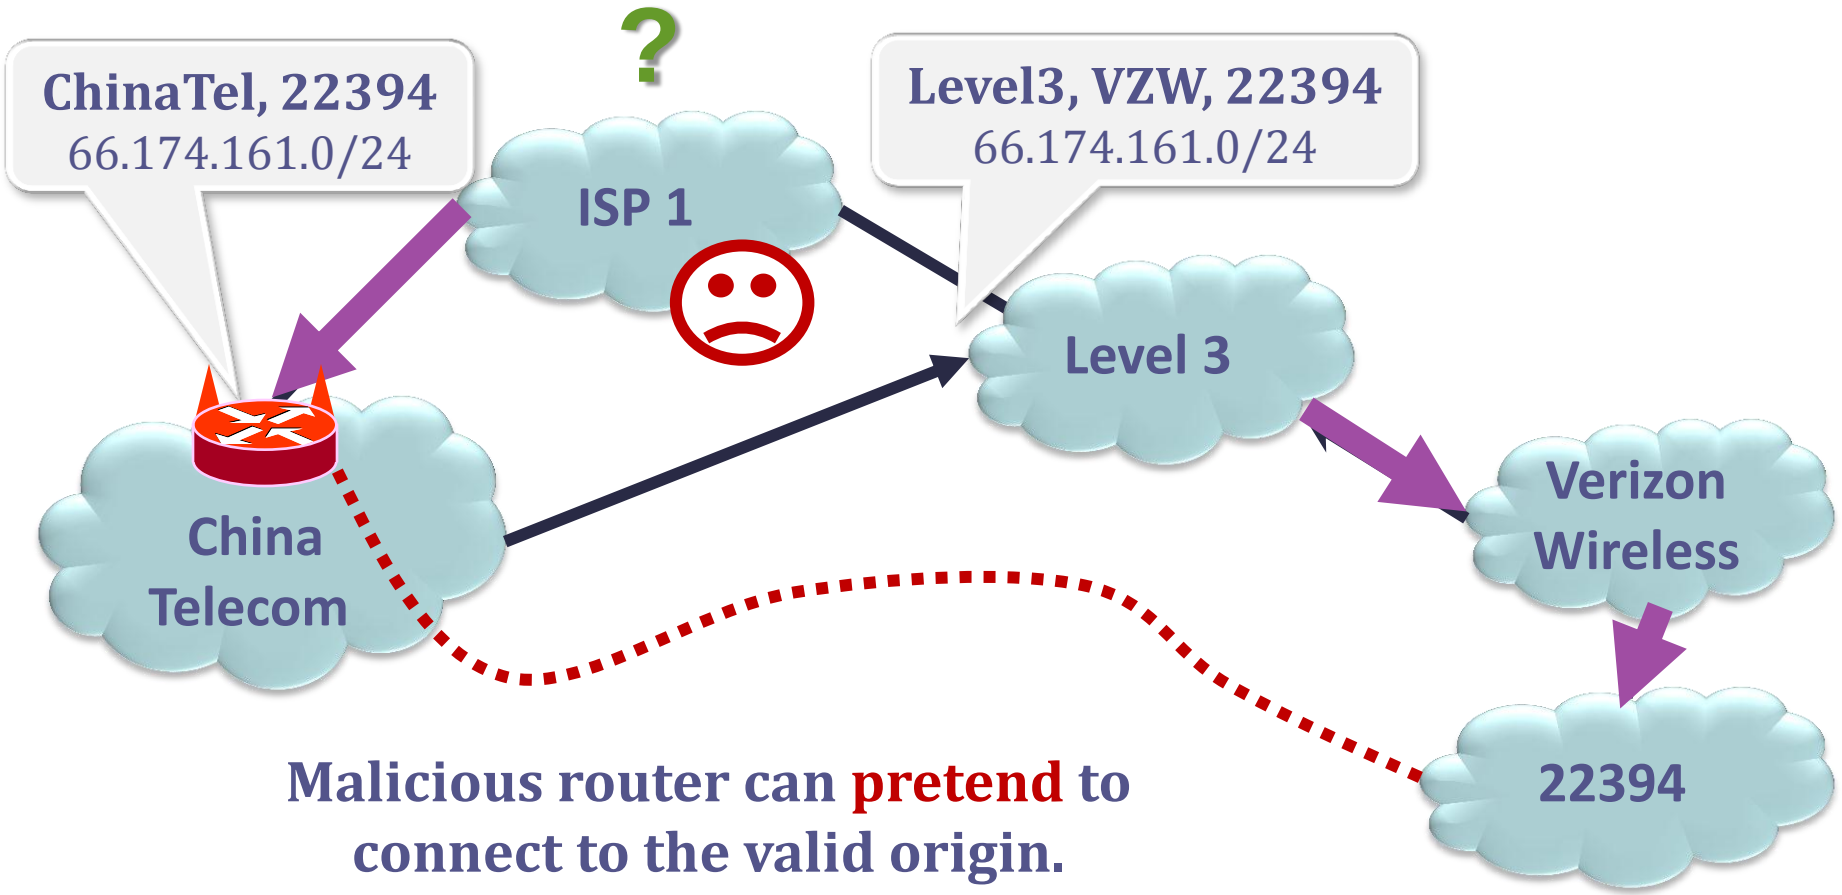

Malicious router can pretend to connect to the valid origin.

To stop this attack, we need S*BGP (e.g. S-BGP/soBGP) (1)

S-BGP [1997]: RPKI + Cannot announce a path that was not announced to you.

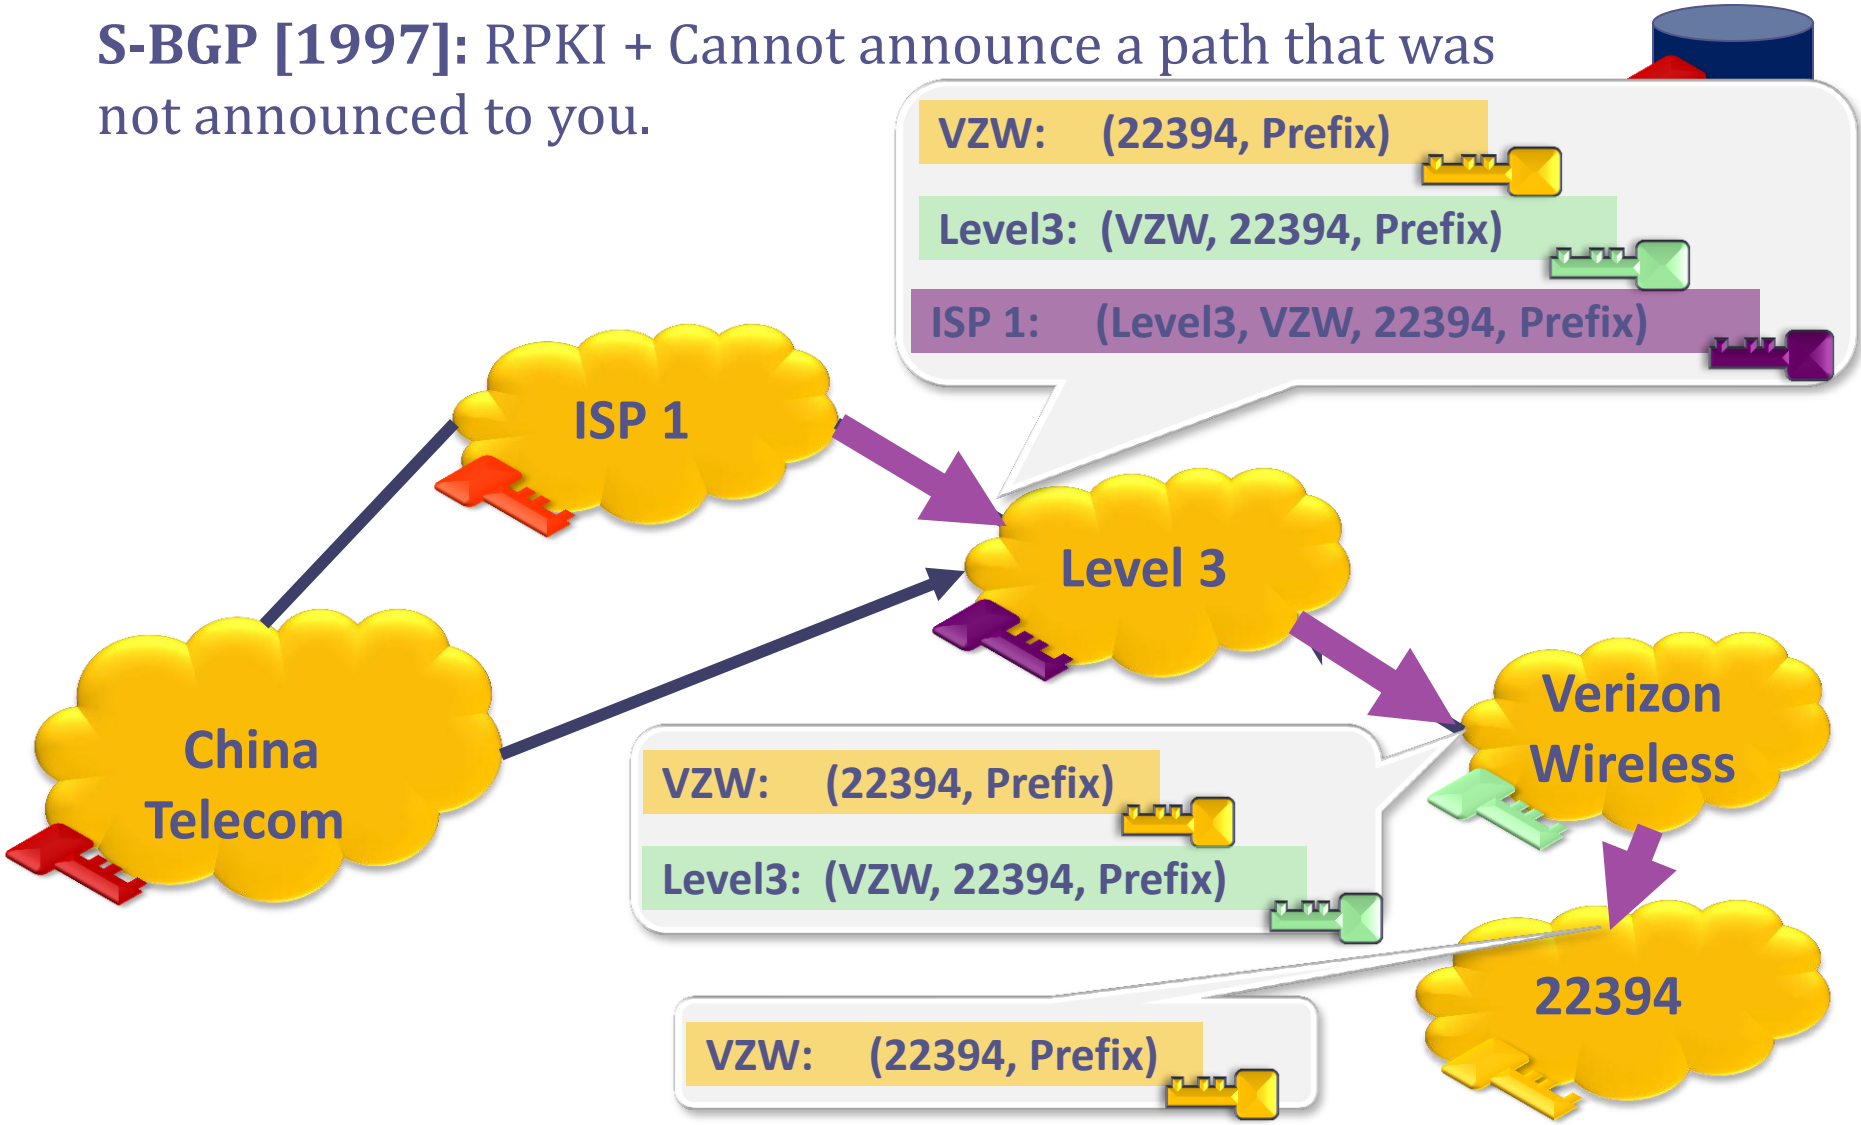

Public Key Signature: Anyone with 22394's public key can validate that the message was sent by 22394.

To stop this attack, we need S*BGP (e.g. S-BGP/soBGP) (2)

S-BGP [1997]: RPKI + Cannot announce a path that was not announced to you.

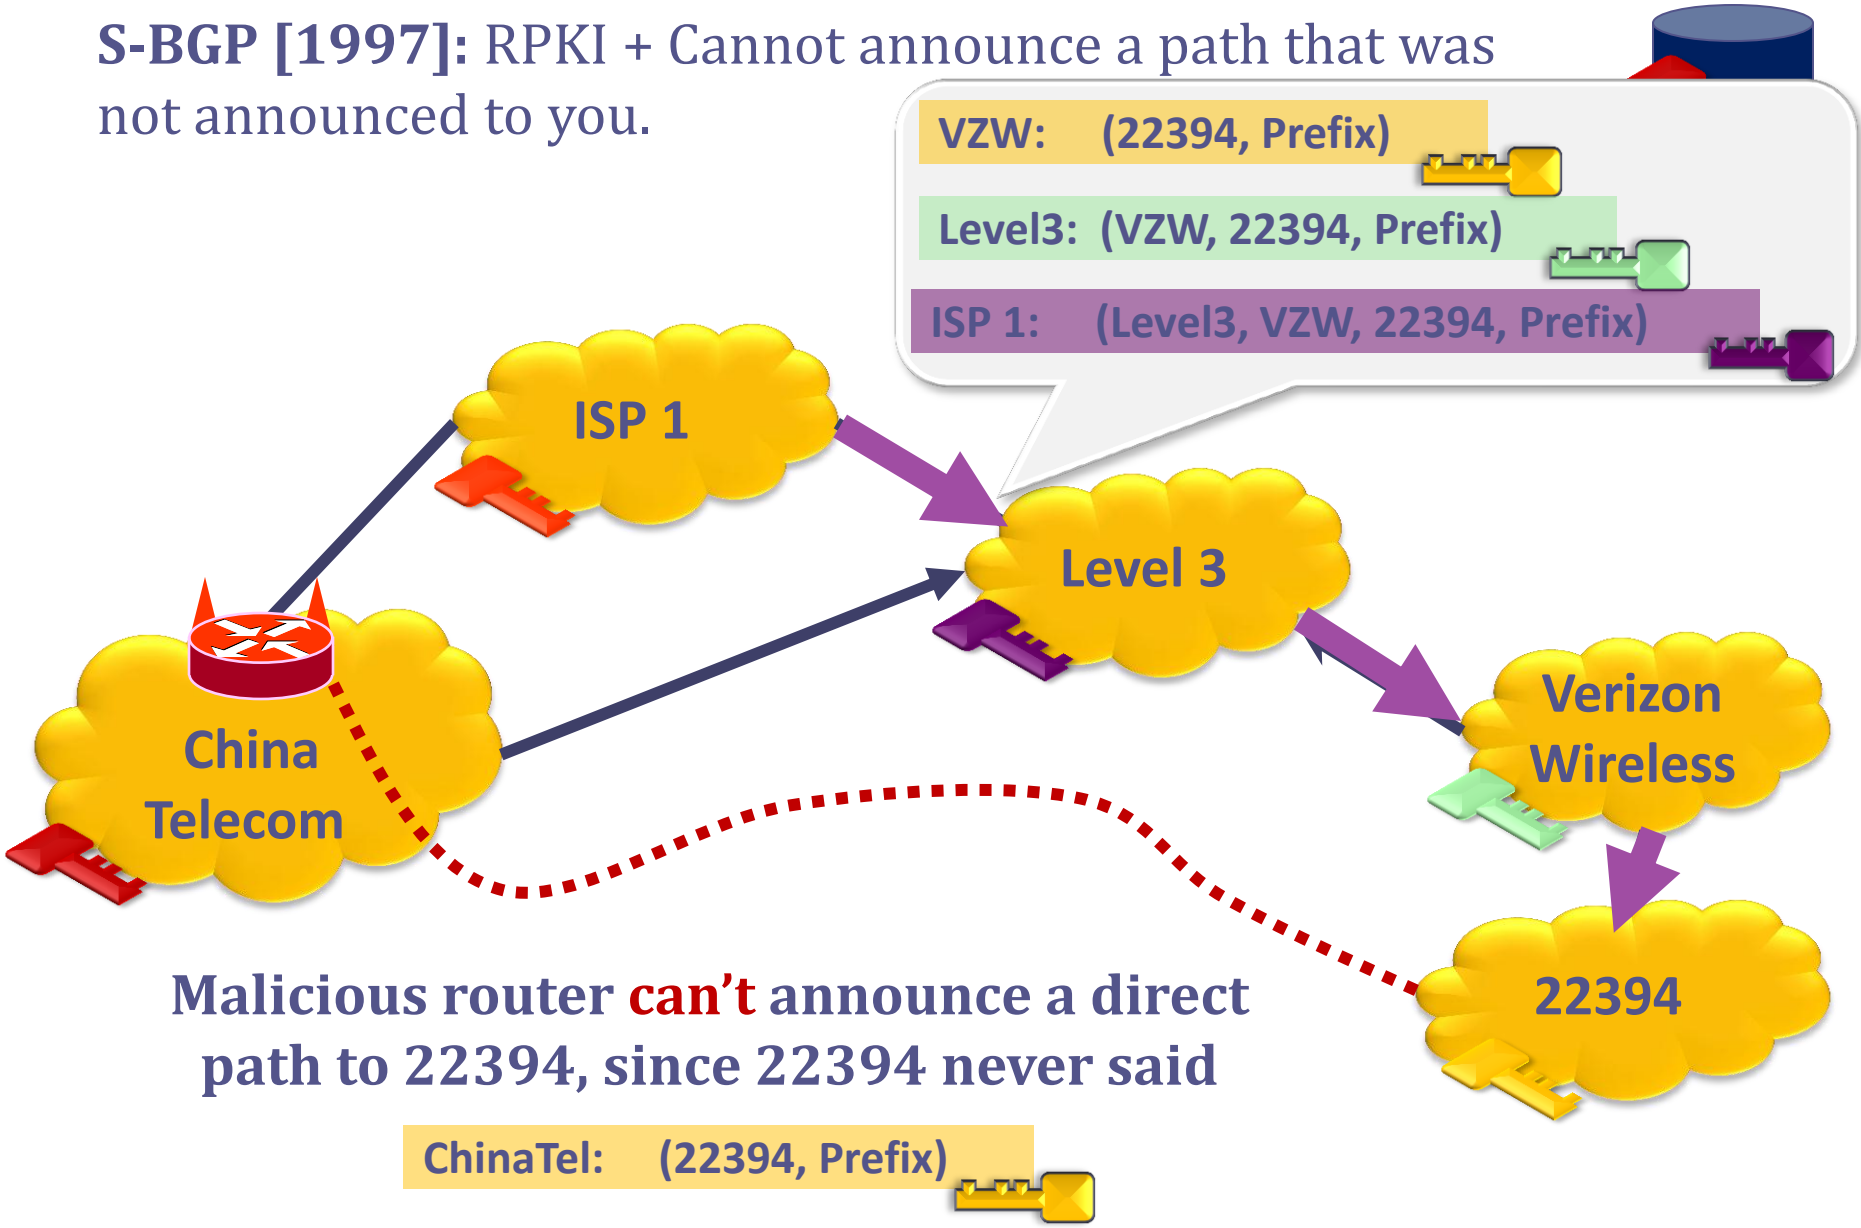

Overview

S*BGP will necessarily go through a transition phase

- How should deployment occur?

Our Goal: Come up with a strategy for S*BGP (S-BGP/soBGP) deployment.

- How governments & standards groups invest resources
- … to create market pressure for S*BGP deployment

We evaluate guidelines via a model & simulations

- Model: ISPs care only about **revenue**, not **security**!
- And run simulations on [UCLA Cyclops+IXP] AS graph data
- Parallelize simulations on a 200-node DryadLINQ cluster

Outline

- Part 1: Background
- Part 2: Our strategy
- Part 3: Evaluating our strategy
  - Model
  - Simulations
- Part 4: Summary and recommendations

How to deploy S*BGP globally?

Pessimistic view:

- No local **economic** incentives; only security incentives
- Like IPv6, but worse, because entire path must be secure

Our view:

- S*BGP has an advantage: **it affects route selection**
- Route selection controls traffic flows
- And an ISP that attracts more customer traffic earns more revenue.

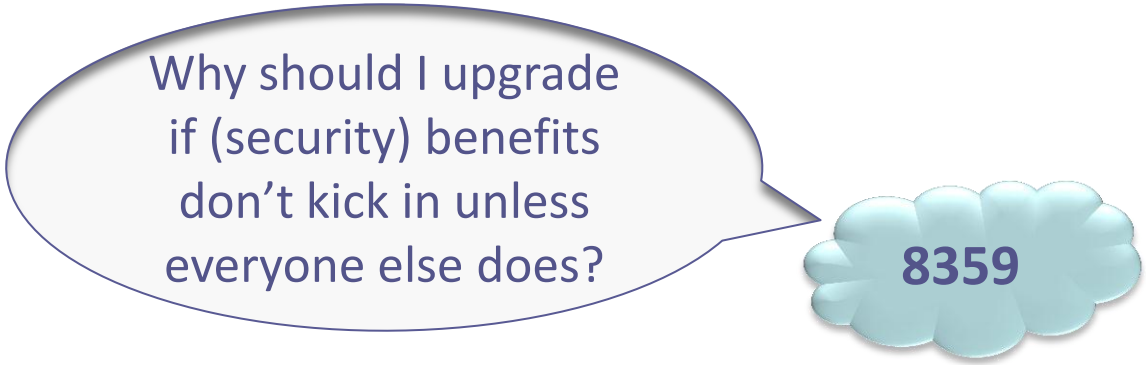

Stubs vs ISPs: Stubs are 85% of the Internet's ASes!

A stub is an AS with no customers.

Stubs shouldn't transit traffic. They only originate their own prefixes.

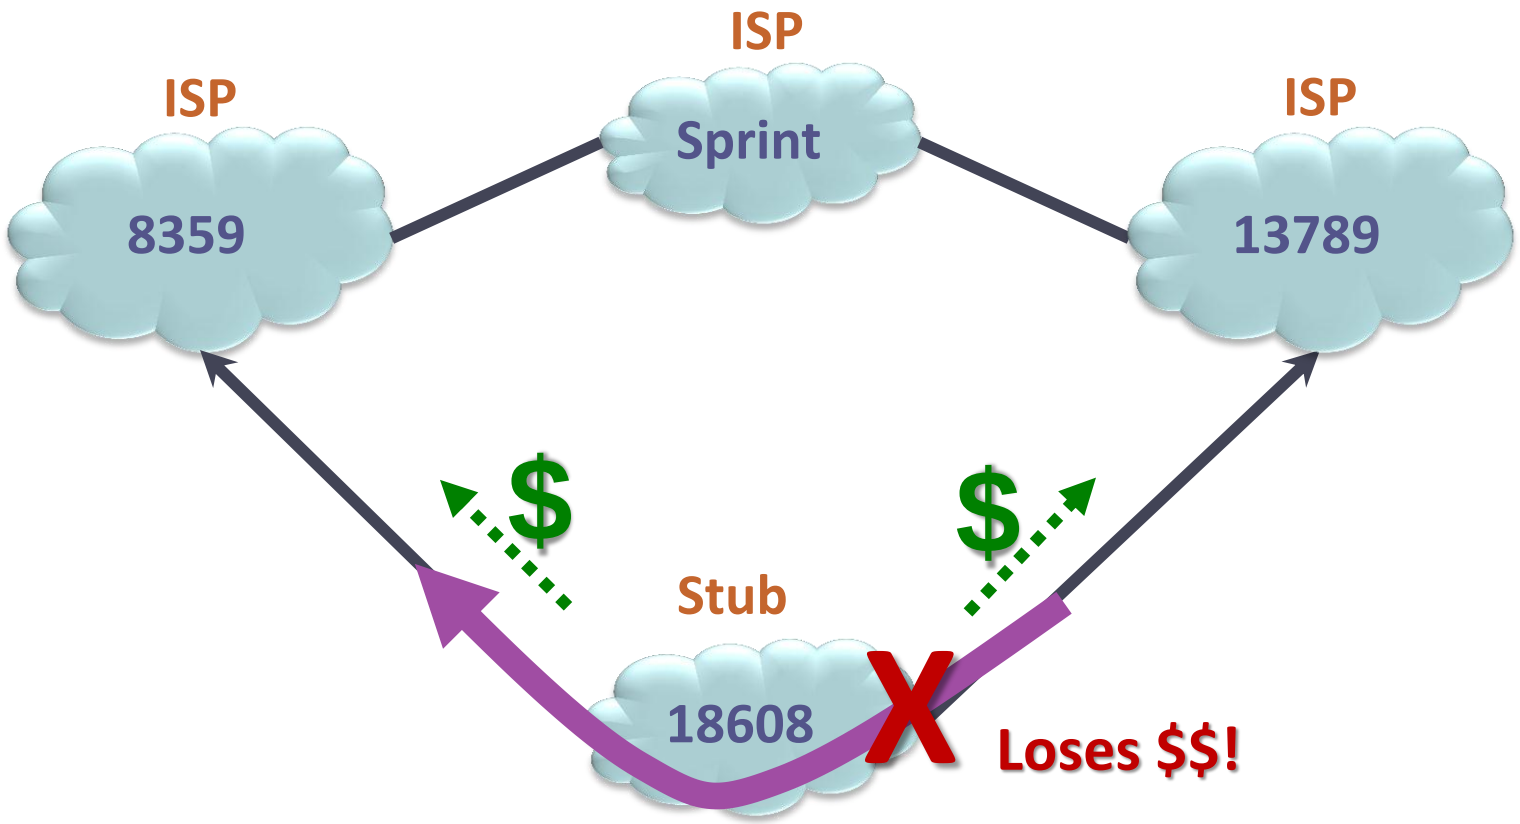

85% of ASes are stubs! **We call the rest (15%) ISPs.**

How can we create market pressure?

Assume that secure ASes *break ties* on secure paths!

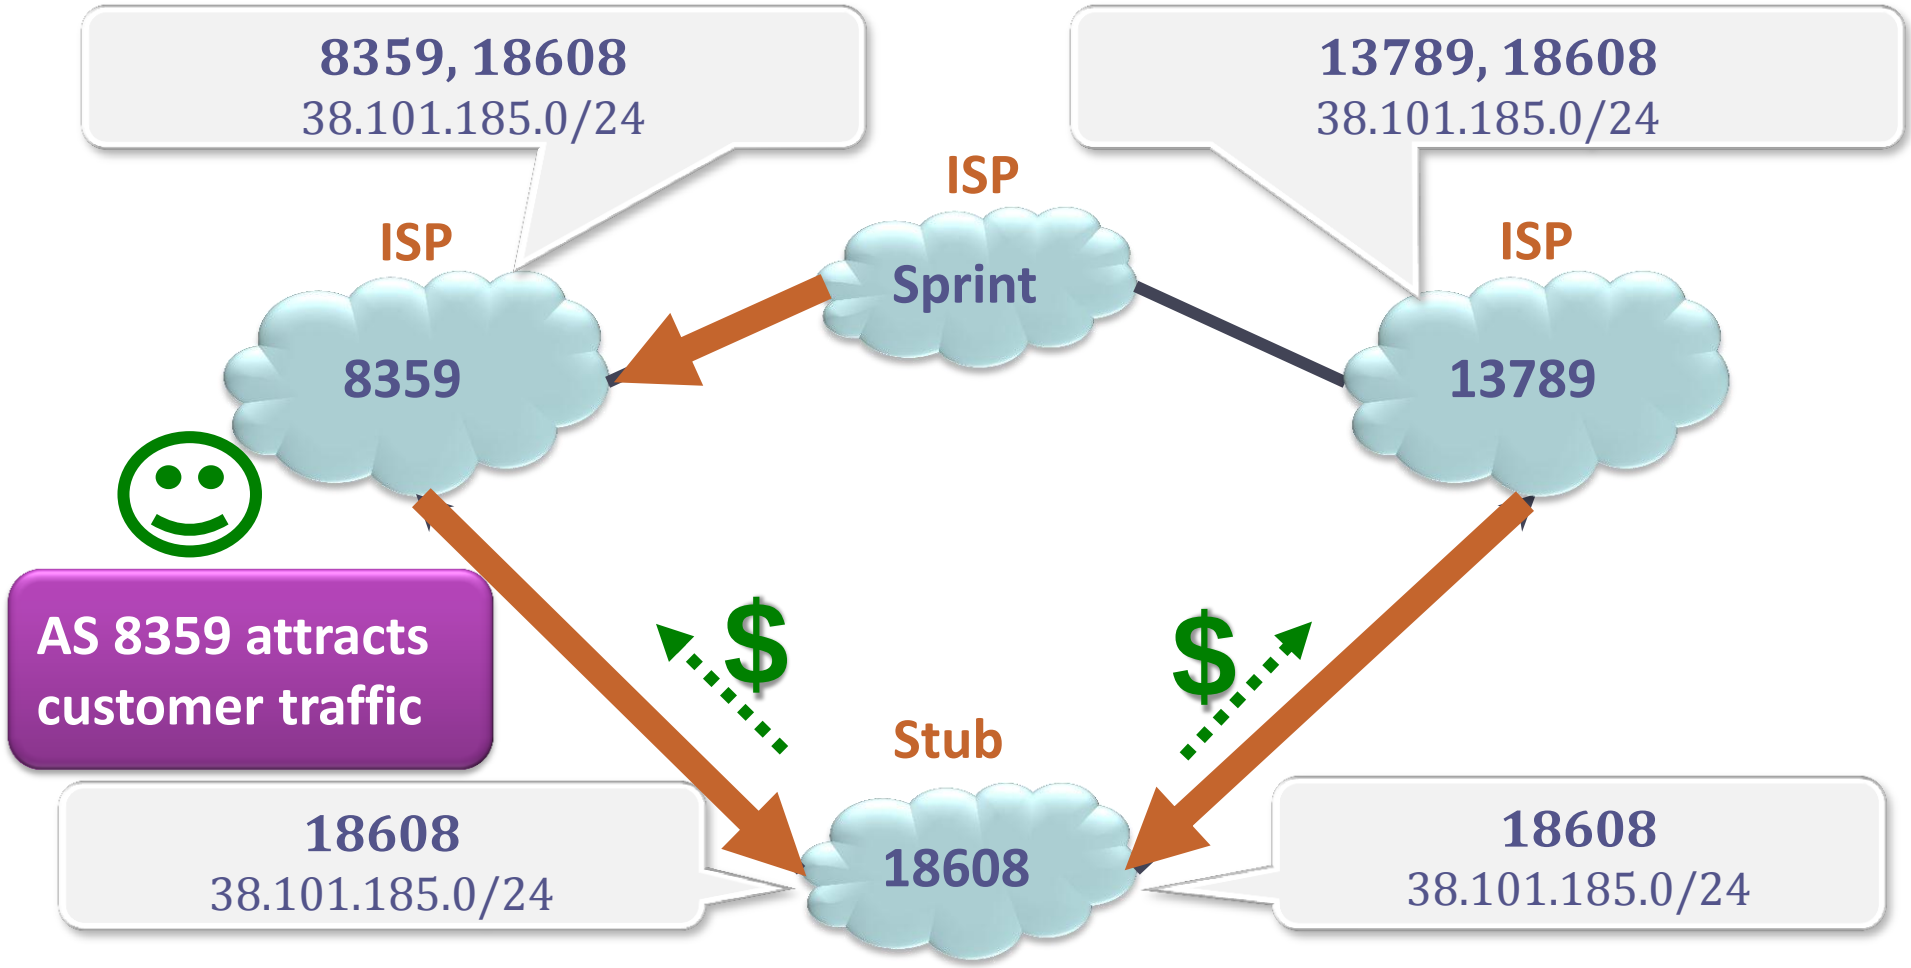

ISPs can use S*BGP to attract customer traffic & thus money

How can we create market pressure?

Assume that secure ASes *break ties* on secure paths!

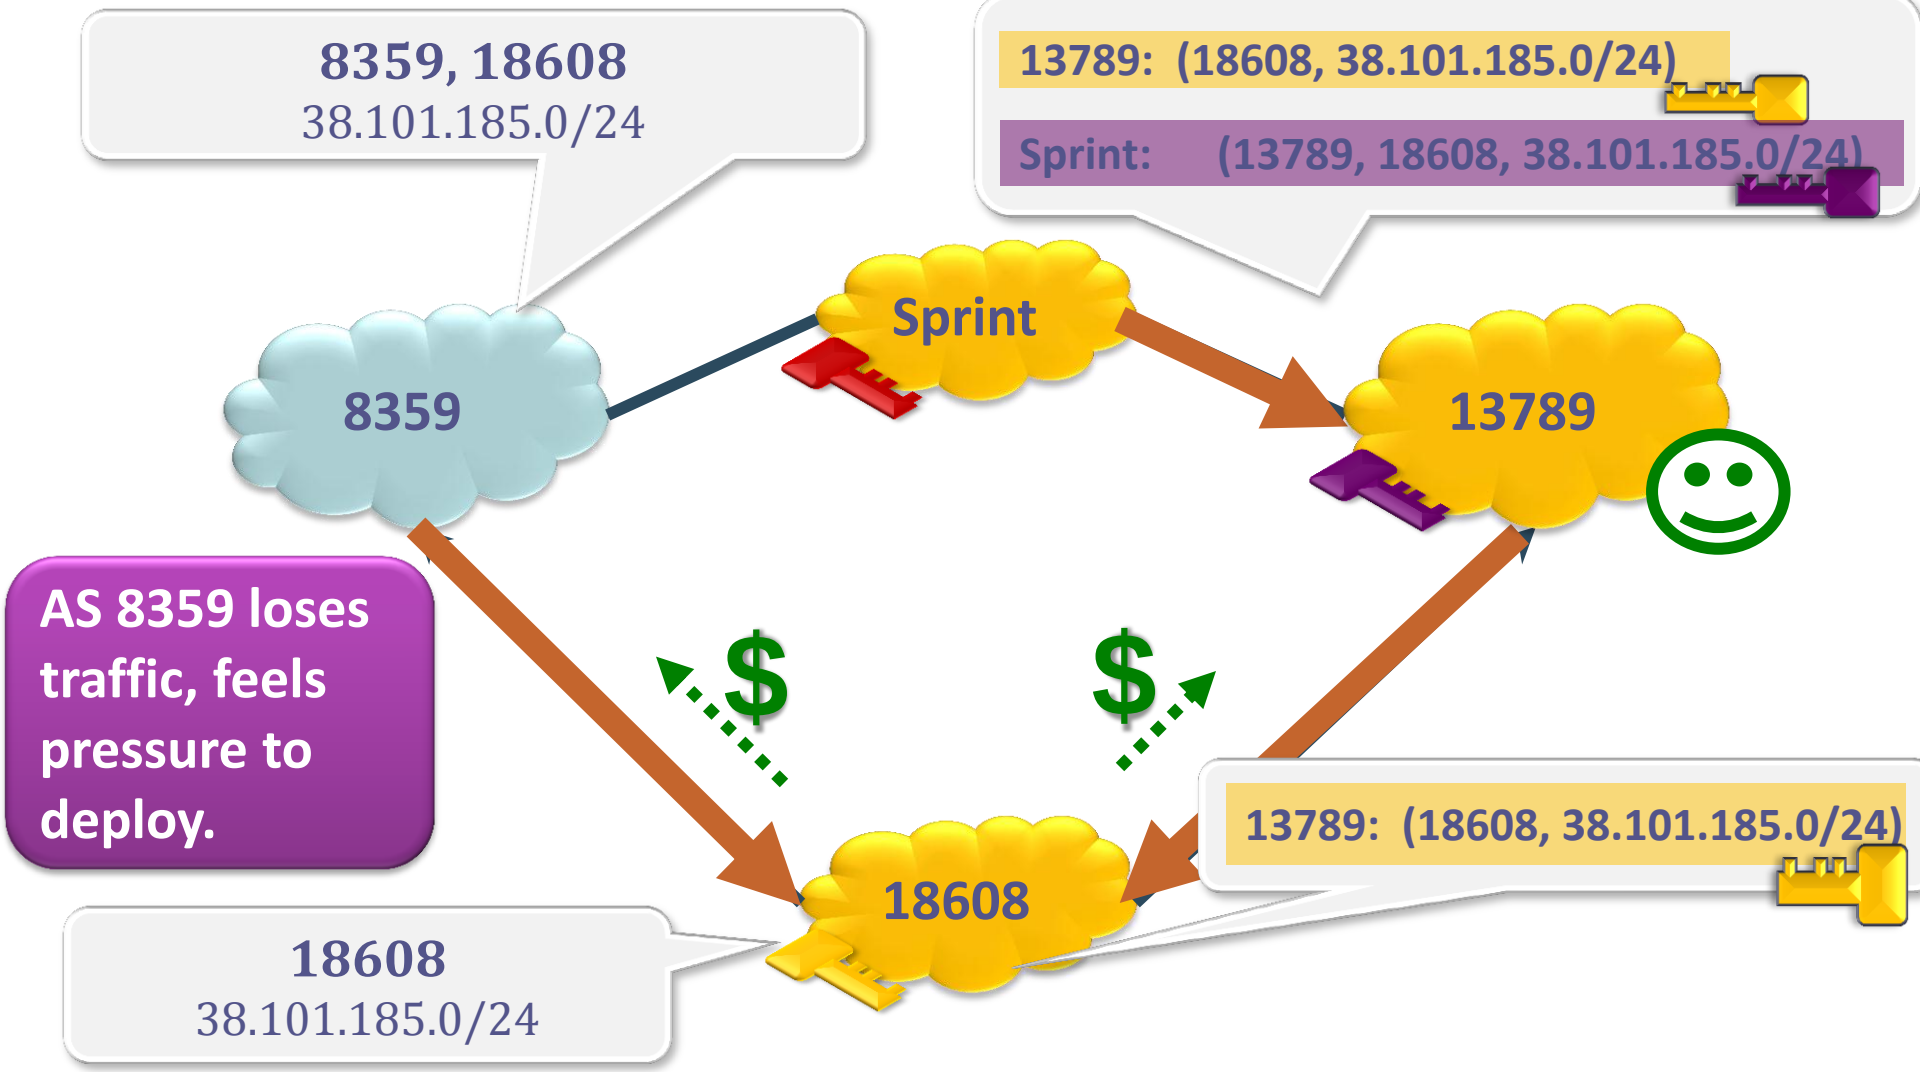

Our Strategy: 3 Guidelines for Deploying S*BGP (1)

1. Secure ASes should break ties in favor of **secure paths**

2. ISPs "help" their **stub** customers deploy **simplex S*BGP.**

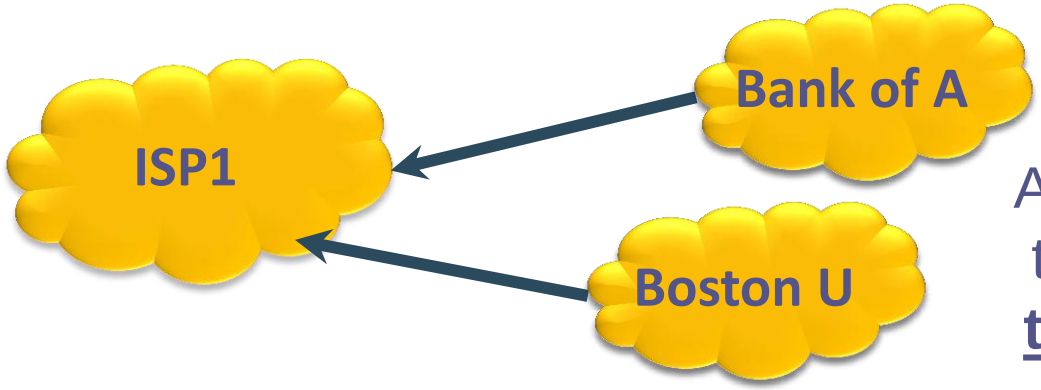

A stub is an AS that **does not transit traffic.**

85% of ASes are stubs!

Simplex S*BGP: `Cheap' S*BGP for Stubs

A stub never transits traffic

- Only announces its own prefixes..
- ...and receives paths from provider
- **Sign but don't verify!**

(rely on provider to validate)

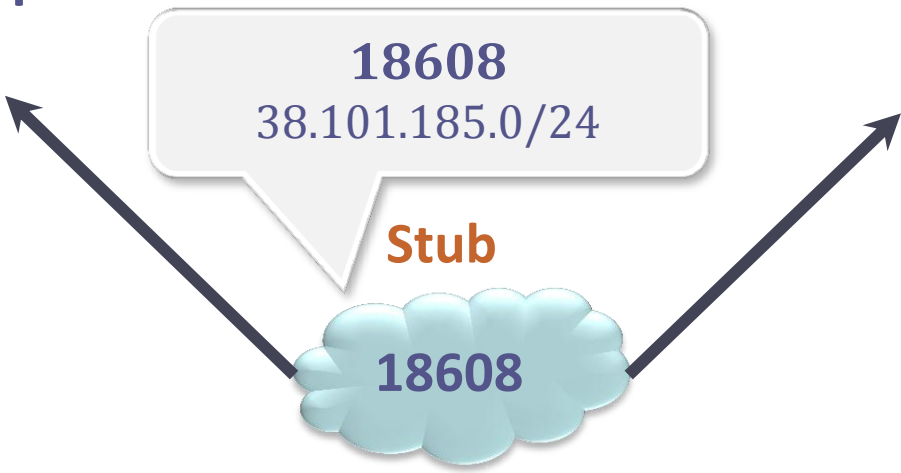

2 options for deploying S*BGP in stubs:

1. Have providers sign for stub customers. (Stubs do nothing)
2. Stubs run **simplex S*BGP.** (Stub only signs, provider validates)
   1. **No hardware upgrade required**
      - Sign for **~1 prefix**, not **~300K prefixes**
      - Use **~1 private key,** not **~36K public keys**
   2. **Security impact is minor (we evaluated this):**
      - Stub vulnerable to attacks by its direct provider.

Our Strategy: 3 Guidelines for Deploying S*BGP (2)

1. Secure ASes should break ties in favor of **secure paths**

2. ISPs "help" their **stub** customers deploy **simplex S*BGP.**

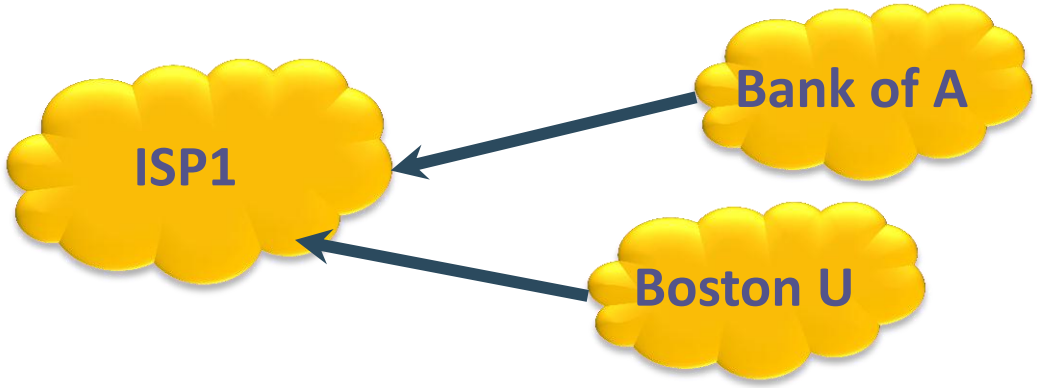

(possibly with some government subsidies)

3. Initially, a few **early adopters** deploy S*BGP (gov't incentives, regulations, altruism, etc).

Outline

- Part 1: Background
- Part 2: Our strategy
- Part 3: Evaluating our strategy
  - Model
  - Simulations
- Part 4: Summary and recommendations

A model of the S*BGP deployment process

- **To start the process:**
  - Early adopter ASes have deployed S*BGP
  - Their stub customers deploy simplex S*BGP

- **Each round:**
  - Compute **utility** for **every insecure** ISP n
  - If ISP n 's utility can increase by more than **θ%** when it deploys S*BGP,
  - Then ISP n decides to **secure itself** & **all its stub** customers

- **Stop when no new ISPs decide to become secure.**

How do we compute utility?

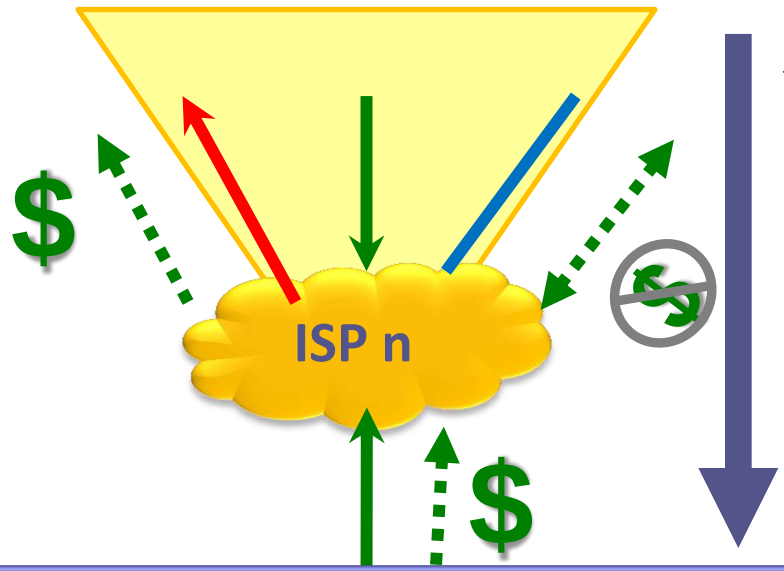

Number of source ASes routing through ISP n **to all customer destinations.**

Important Note: ISP utility does not depend on security.

To determine routing, we run simulations on the [UCLA Cyclops] AS graph with these routing policies:

BGP Routing Policy Model:

1. Prefer customer paths
   over peer paths
   over provider paths
2. Prefer shorter paths
3. **If secure, prefer secure paths**
4. Arbitrary tiebreak

Outline

- Part 1: Background
- Part 2: Our strategy
- Part 3: Evaluating our strategy
  - Model
  - Simulations
- Part 4: Summary and recommendations

Case Study of S*BGP deployment

Ten early adopters:

- **Five Tier 1s:**
  - Sprint (AS 1239)
  - Verizon (AS 701)
  - AT&T (AS 7018)
  - Level 3 (AS 3356)
  - Cogent (AS 174)
- **Five Popular Content Providers**
  - Google (AS 15169)
  - Microsoft (AS 8075)
  - Facebook (AS 32934)
  - Akamai (AS 22822)
  - Limelight (AS 20940)

- The five content providers source **10%** of Internet traffic
- Stubs break ties in favor of secure paths
- Threshold **$\theta = 5\%$.**

This leads to 85% of ASes deploying S*BGP (65% of ISPs)

Simulation: Market pressure drives deployment (1)

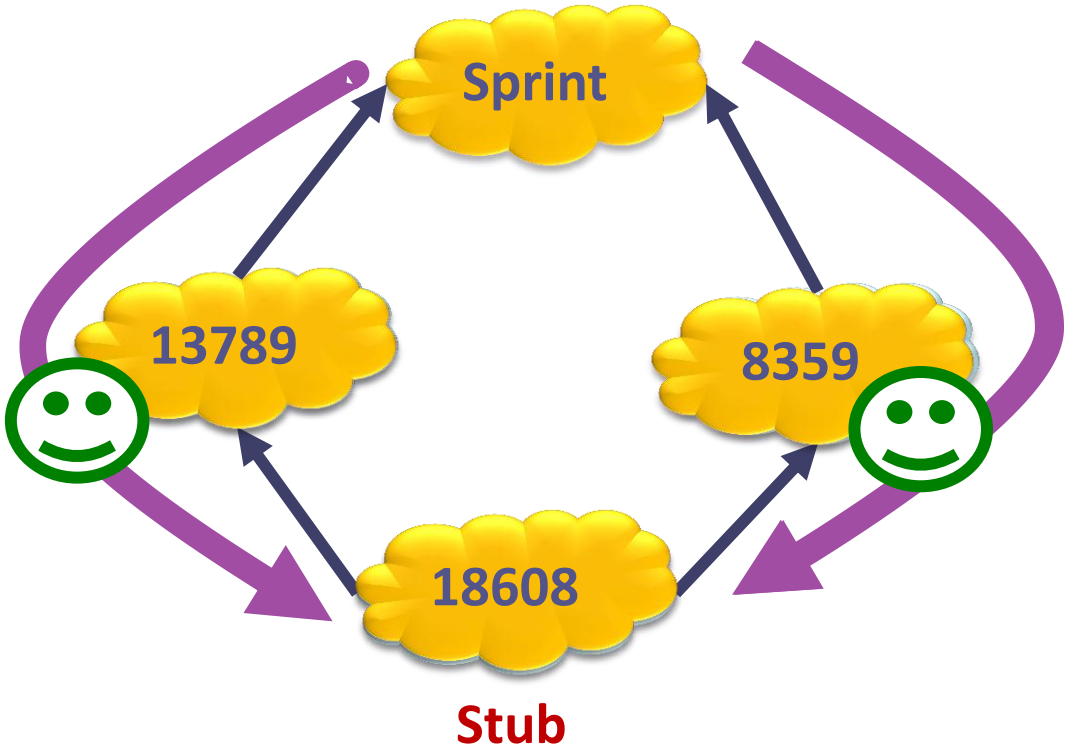

Simulation: Market pressure drives deployment (2)

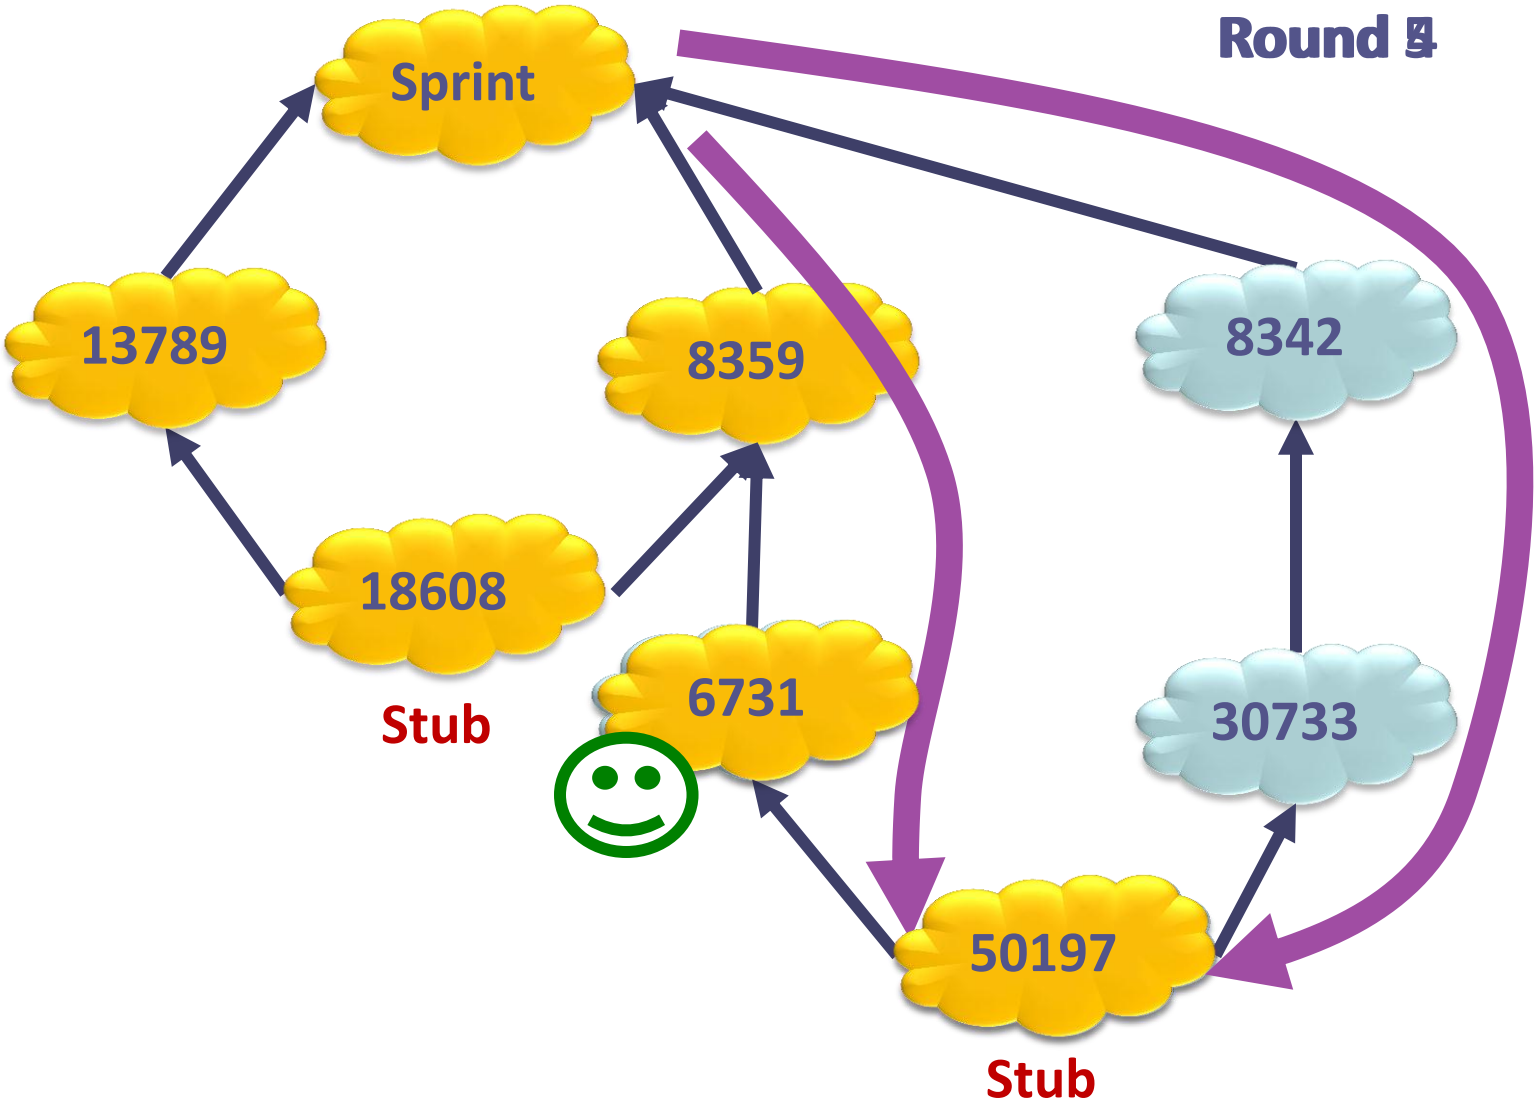

Simulation: Market pressure drives deployment (3)

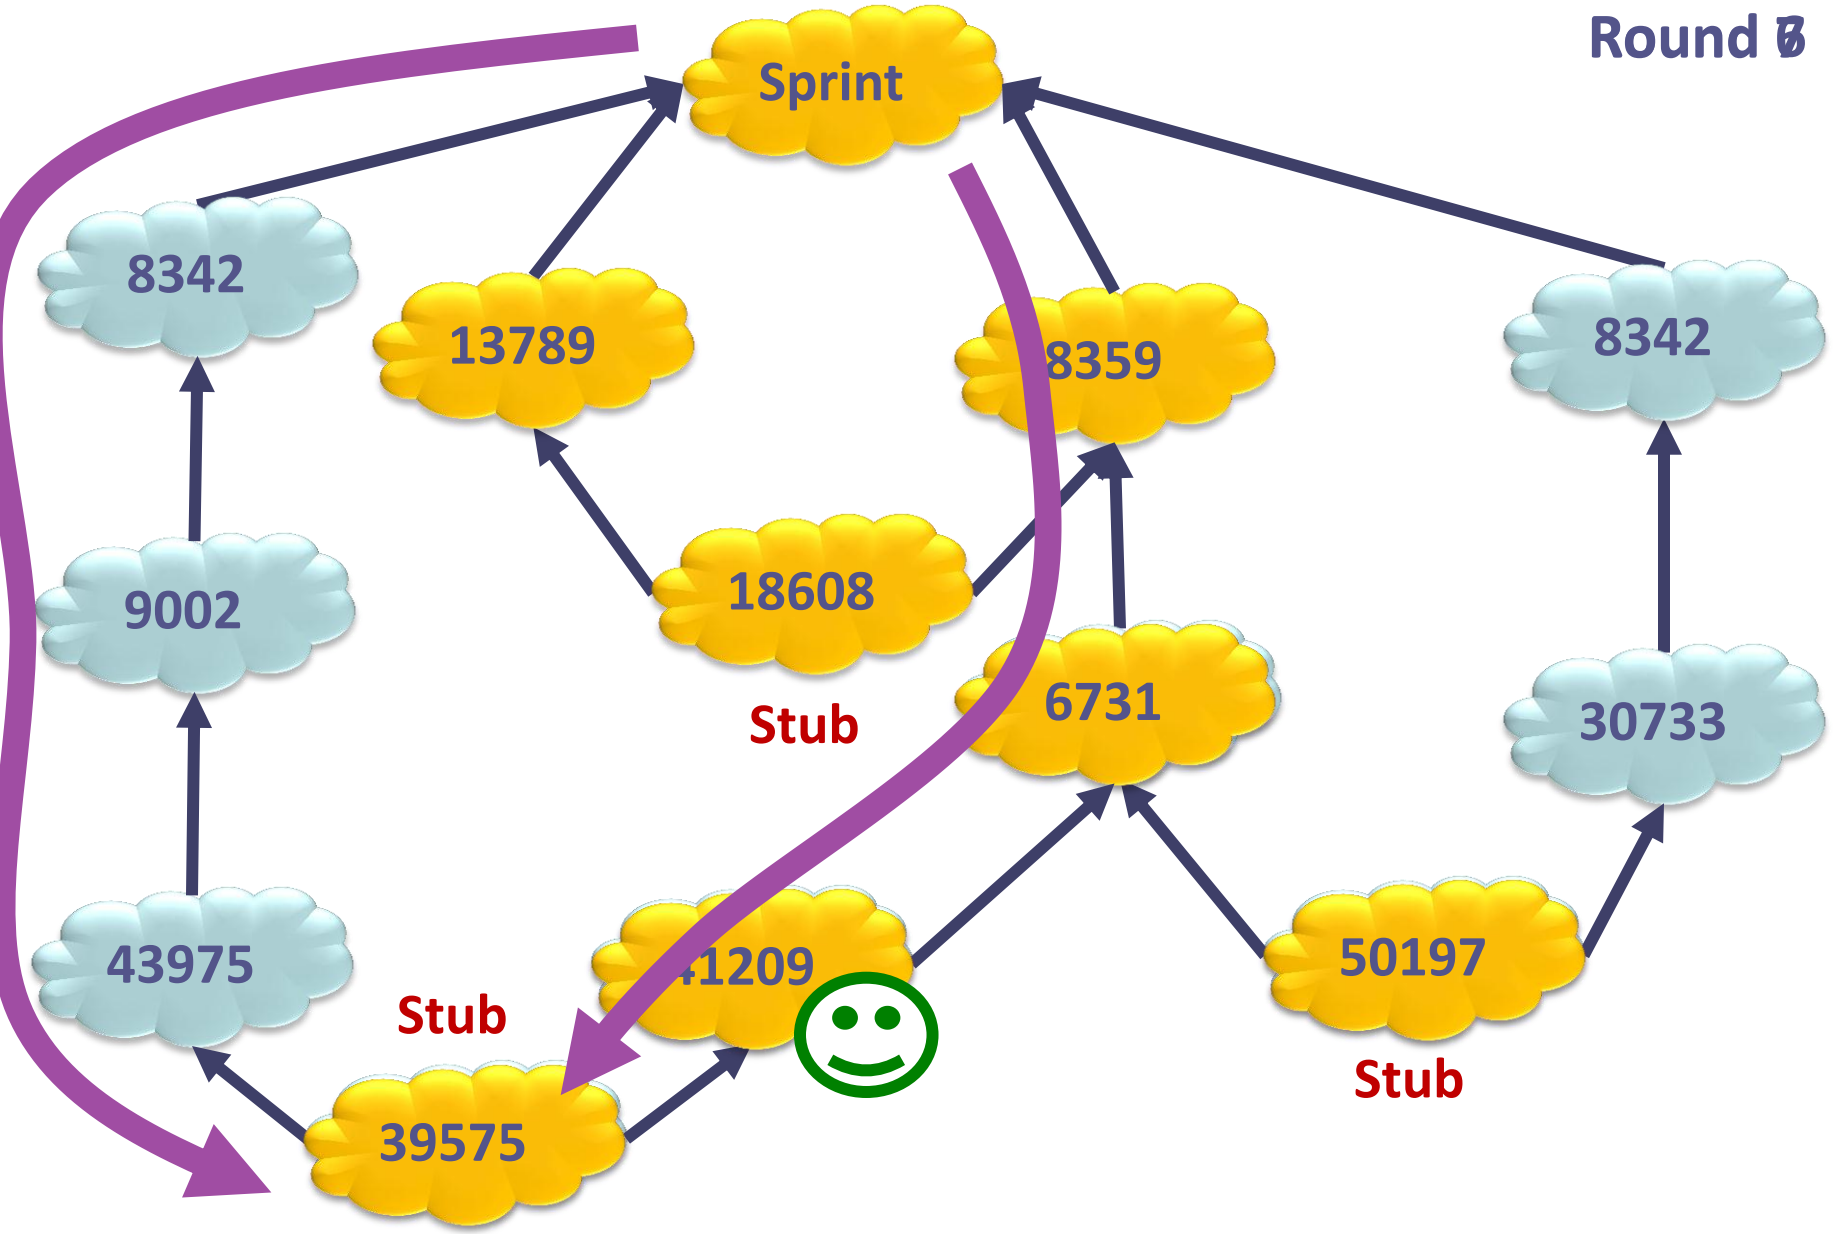

So who should be the early adopters?

Theorem: Finding the optimal set of early adopters is NP-hard. Approximating this within a constant factor is also NP-hard.

So who should be the early adopters?

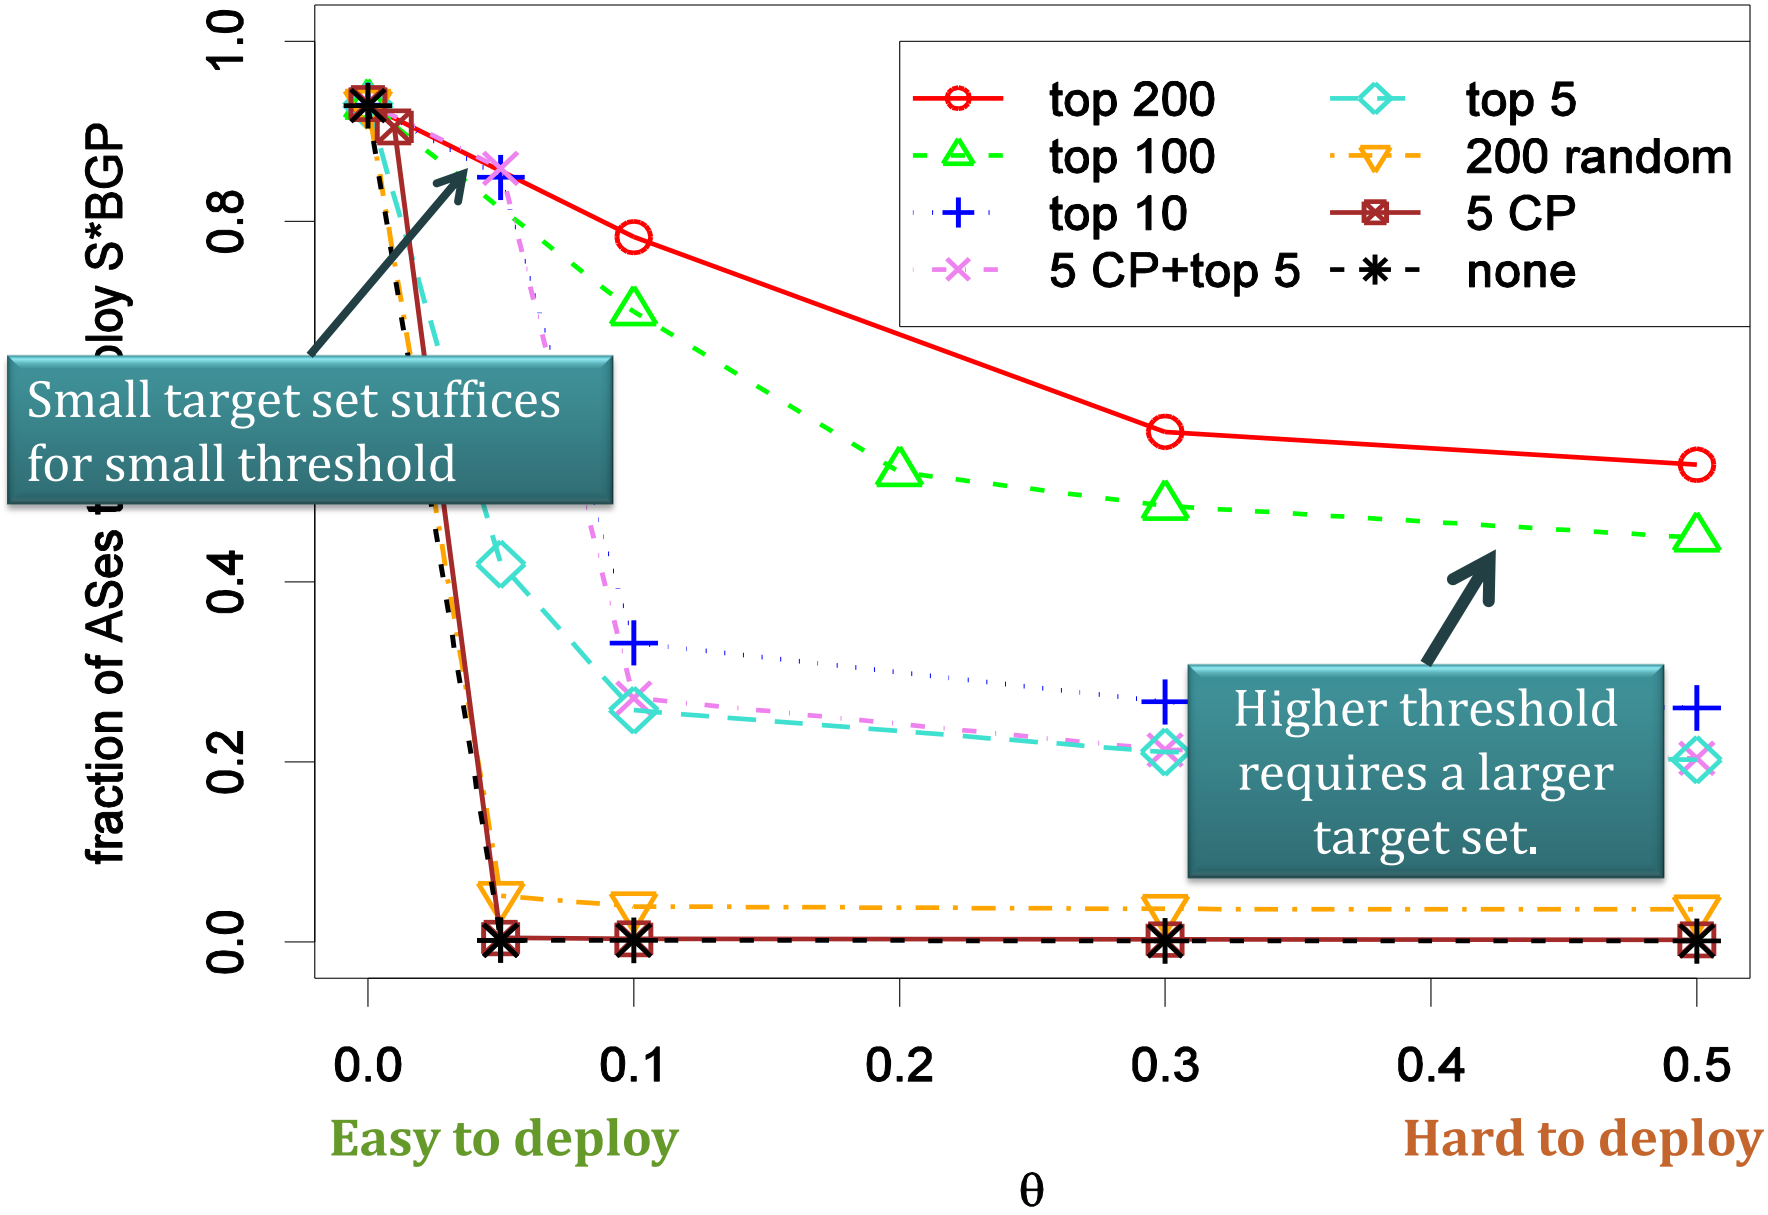

Outline

- Part 1: Background
- Part 2: Our strategy
- Part 3: Evaluating our strategy
  - Model
  - Simulations
- Part 4: Summary and recommendations

Summary and Recommendations

How to create a market for S*BGP deployment?

1. Many secure destinations via simplex S*BGP.
2. Market pressure via S*BGP influence on route selection.

Where should government incentives and regulation go?

1. Focus on early adopters; Tier 1s, maybe content providers
2. Subsidize ISPs to upgrade stubs to simplex S*BGP

Other challenges and future work :

- ISPs can have incentives to turn off S*BGP
- BGP and S*BGP will coexist in the long run
- ISPs need tools to predict S*BGP impact on traffic

Contact: phillipa@cs.toronto.edu

http://www.cs.toronto.edu/~phillipa/sbgpTrans.html

Thanks to Microsoft Research SVC and New England for supporting us with DryadLINQ.

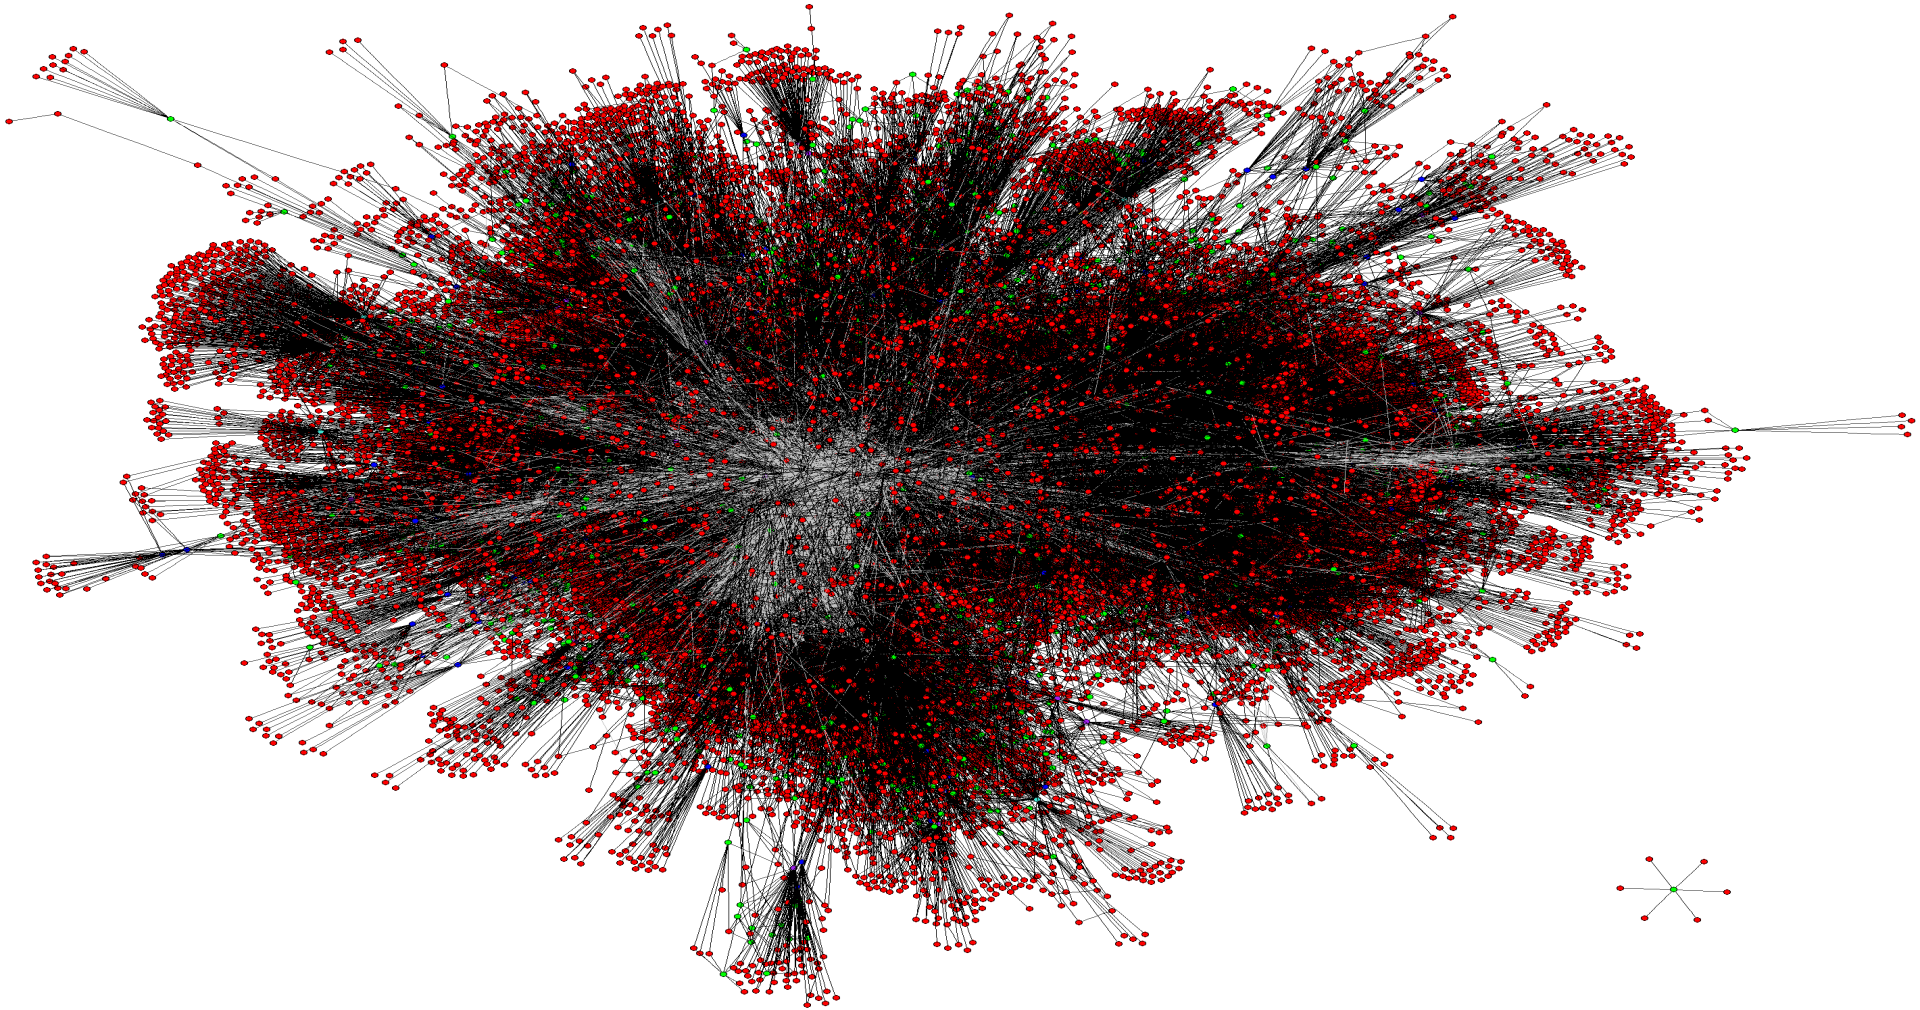

Data Sources for ChinaTel Incident of April 2010

- **Example topology derived from Routeviews messages observed at the LINX Routeviews monitor on April 8 2010**
  - BGP announcements & topology was simplified to remove prepending
  - We anonymized the large ISP in the Figure.
  - Actual announcements at the large ISP were:
  - From faulty ChinaTel router: **"4134 23724 23724 for 66.174.161.0/24"**
  - From Level 3: **"3356 6167 22394 22394 for 66.174.161.0/24"**

- **Traffic interception was observed by Renesys blog**
  - http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml
  - We don't have data on the exact prefixes for which this happened.

- **AS relationships: inferred by UCLA Cyclops**