BOSTON UNIVERSITY

NANOG 52
Denver, CO
June 14, 2011

A Strategy for Transitioning to BGP Security

Sharon Goldberg
Boston University

Phillipa Gill
University of Toronto

Michael Schapira
Princeton University

This work supported by NSF Trustworthy computing grant S-1017907 and a gift from Cisco.

Incentives for BGP Security

What happens after we deploy RPKI? Are we done?

- **NO!** Many attacks on BGP even with RPKI **(See my NANOG'49 talk)**
- Also need path validation with **S*BGP (e.g, BGPsec / soBGP)**
- What are the **incentives** to deploy path validation?

The pessimistic view:

- Why should I bother deploying **S*BGP** in my network?
- No security benefits until many other ASes deploy.
- Worse yet, I can't make money from it.

Our view:

- Calm down. Things aren't so bad.
- You **can** use S*BGP to make money
- …by attracting customers to your ISP.

Overview

Goal of this work:

- We want to engineer the **S*BGP** deployment process
- … so ISPs can make money after they deploy **S*BGP.**
- And we end up with global **S*BGP** deployment

We present & evaluate guidelines for S*BGP deployment.

- Evaluate: model & simulation on **[Cyclops UCLA]** AS graph data
- This talk show results directly from our simulations
- **Caveat**: We **do not predict** how S*BGP deployment will go.
- Our goal is to understand key issues affecting deployment.

Talk Organization

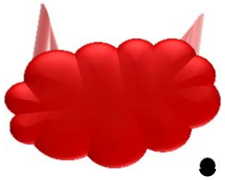

Background:

- BGP, attacks and defenses like **RPKI** & **BGPsec**

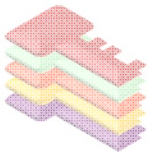

A Strategy for S*BGP deployment

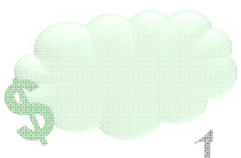

Evaluating our strategy

1. Model
2. Simulation results on **[UCLA Cyclops]** AS graph data

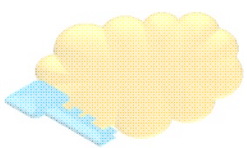

Conclusions and recommendations

Traffic Attraction & Interception Attacks

An interesting incident from April 8, 2010

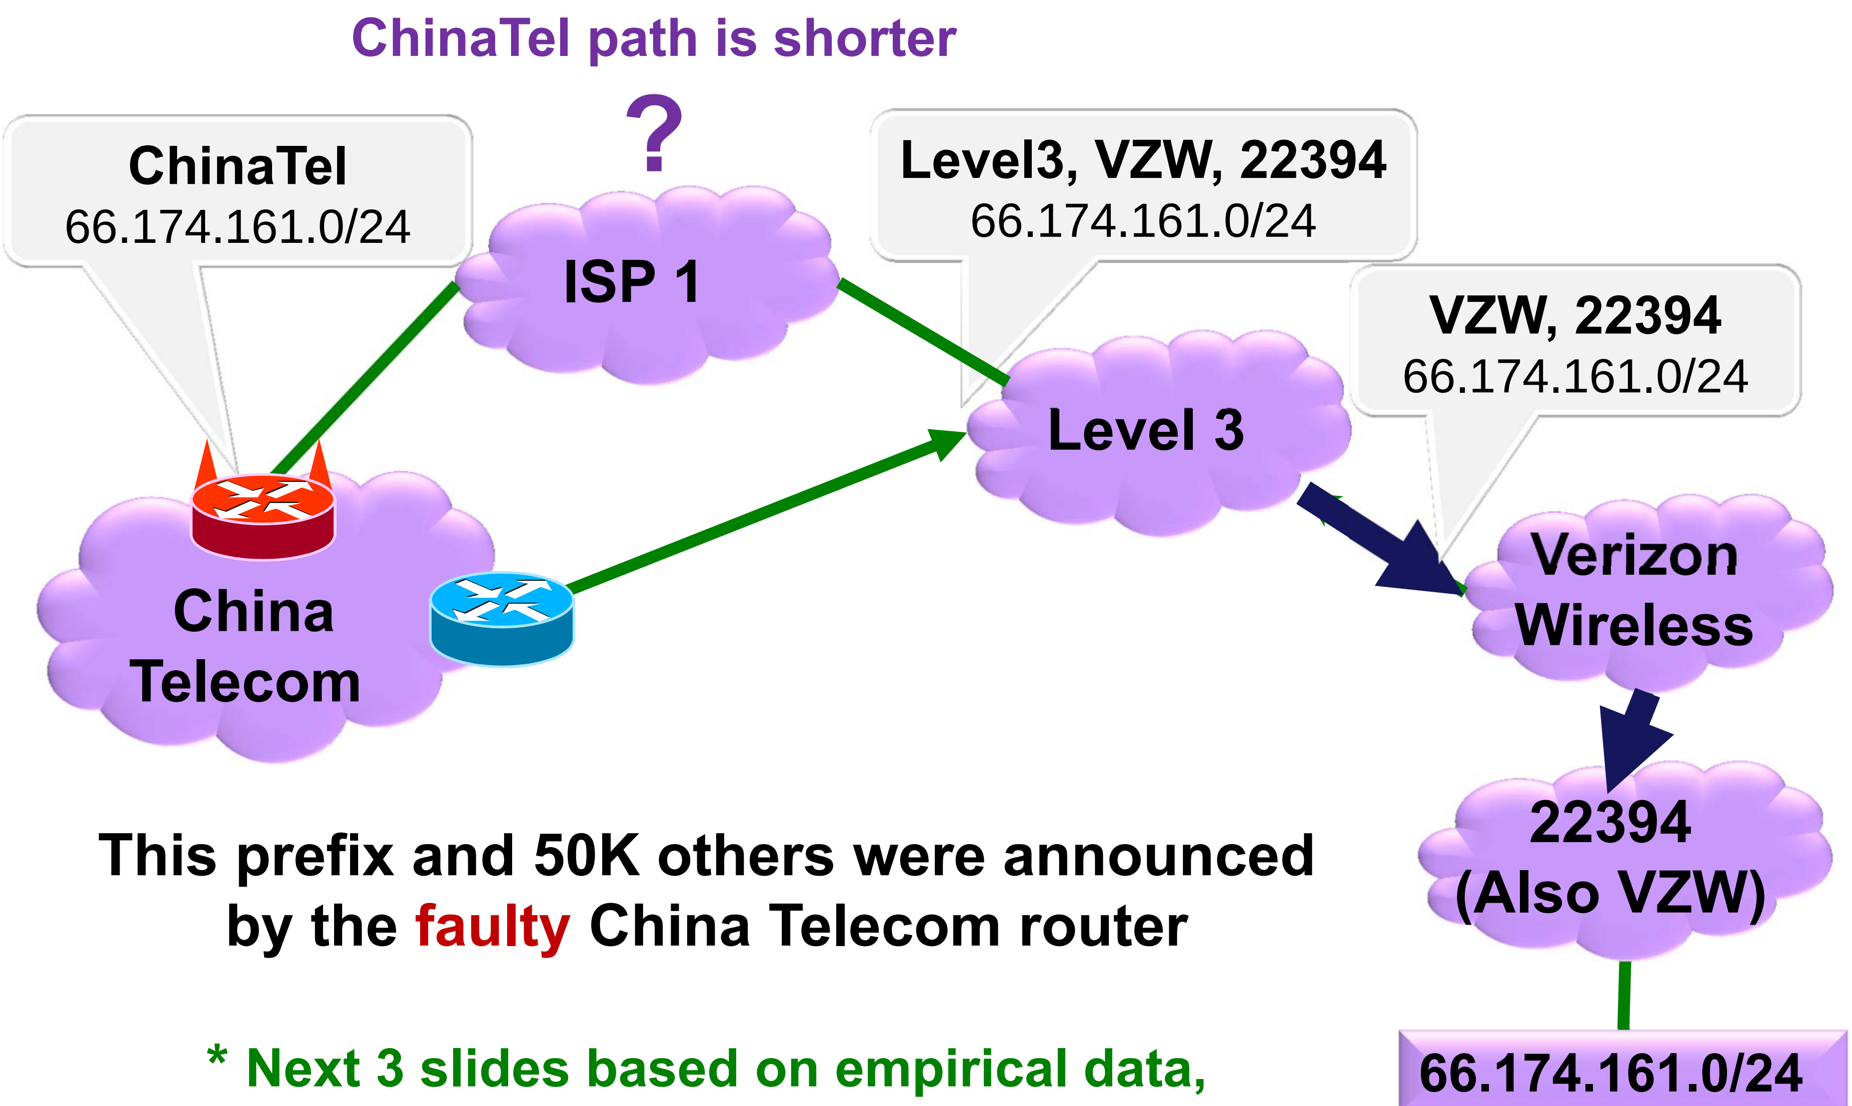

This prefix and 50K others were announced by the faulty China Telecom router

* Next 3 slides based on empirical data, see slide 50 for details on data sources.

Traffic Attraction & Interception Attacks

An interesting incident from April 8, 2010

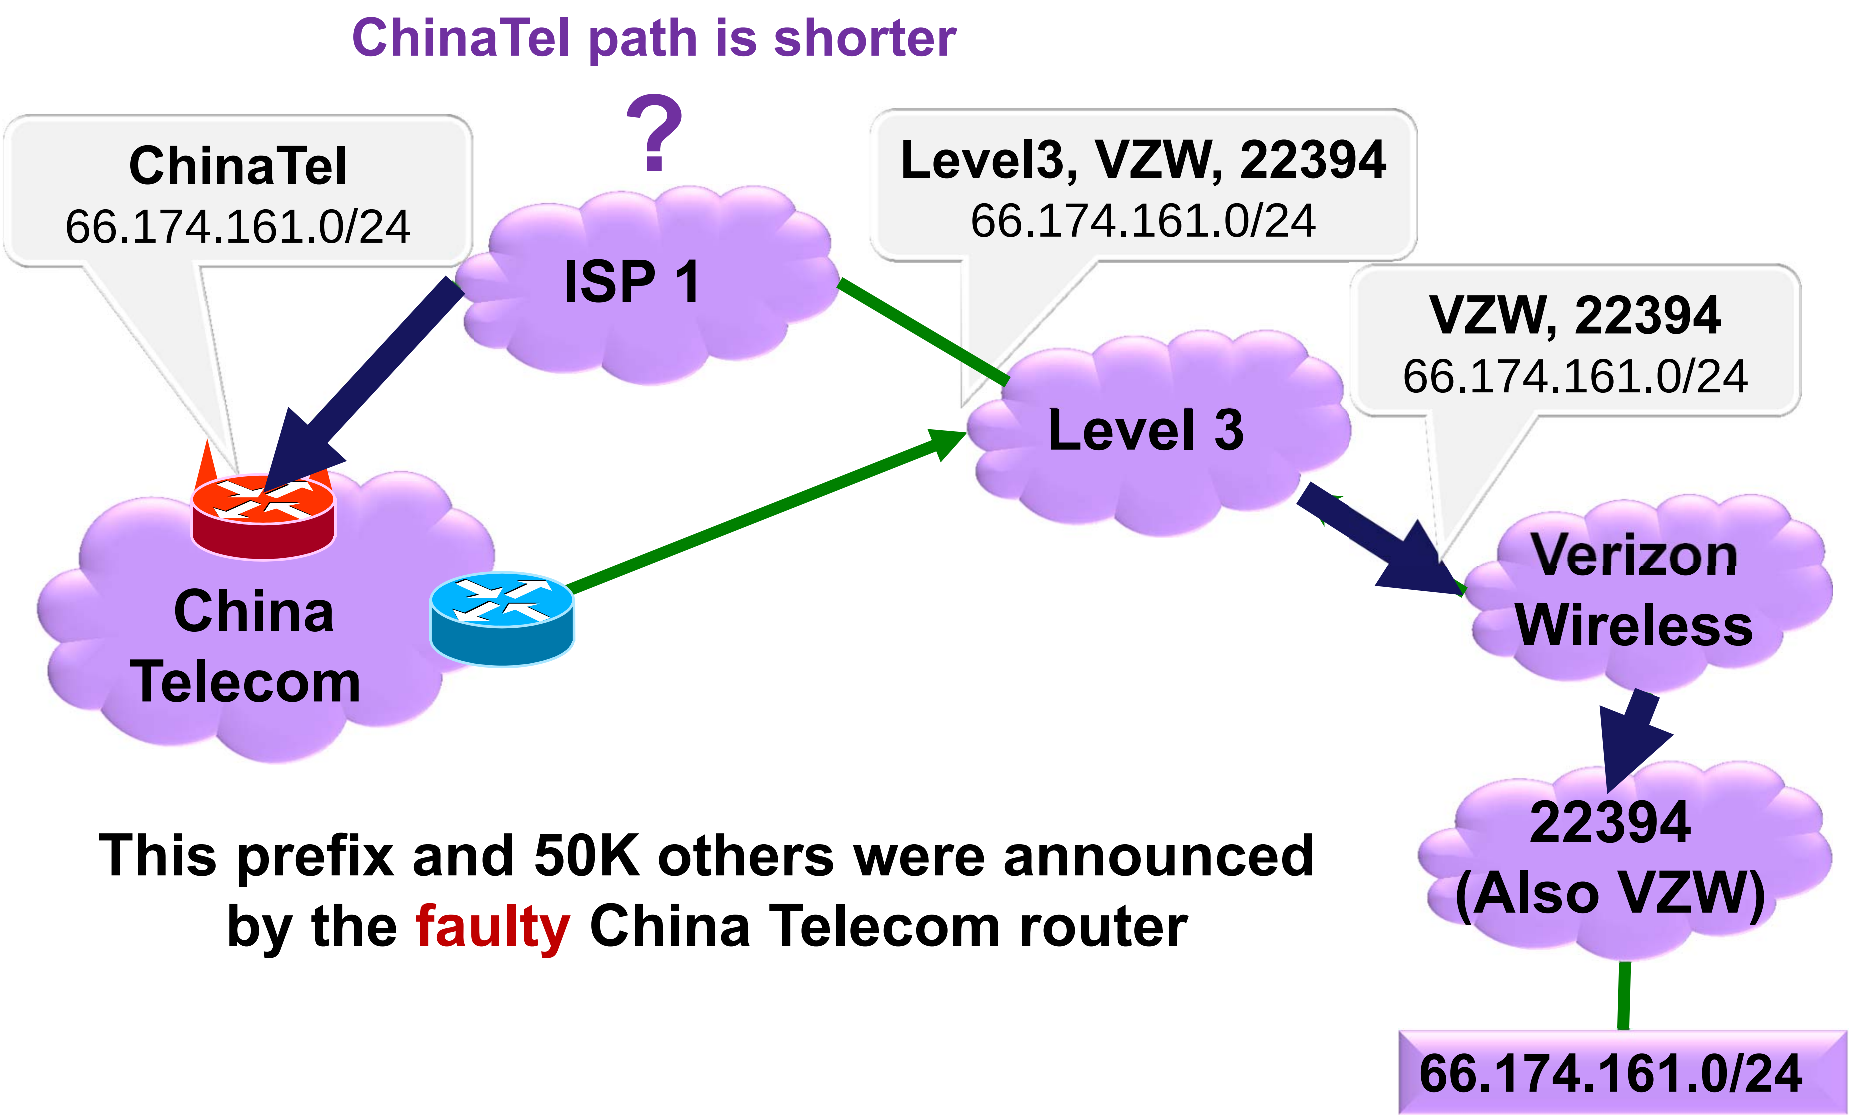

This prefix and 50K others were announced by the faulty China Telecom router

Traffic Attraction & Interception Attacks

An interesting incident from April 8, 2010

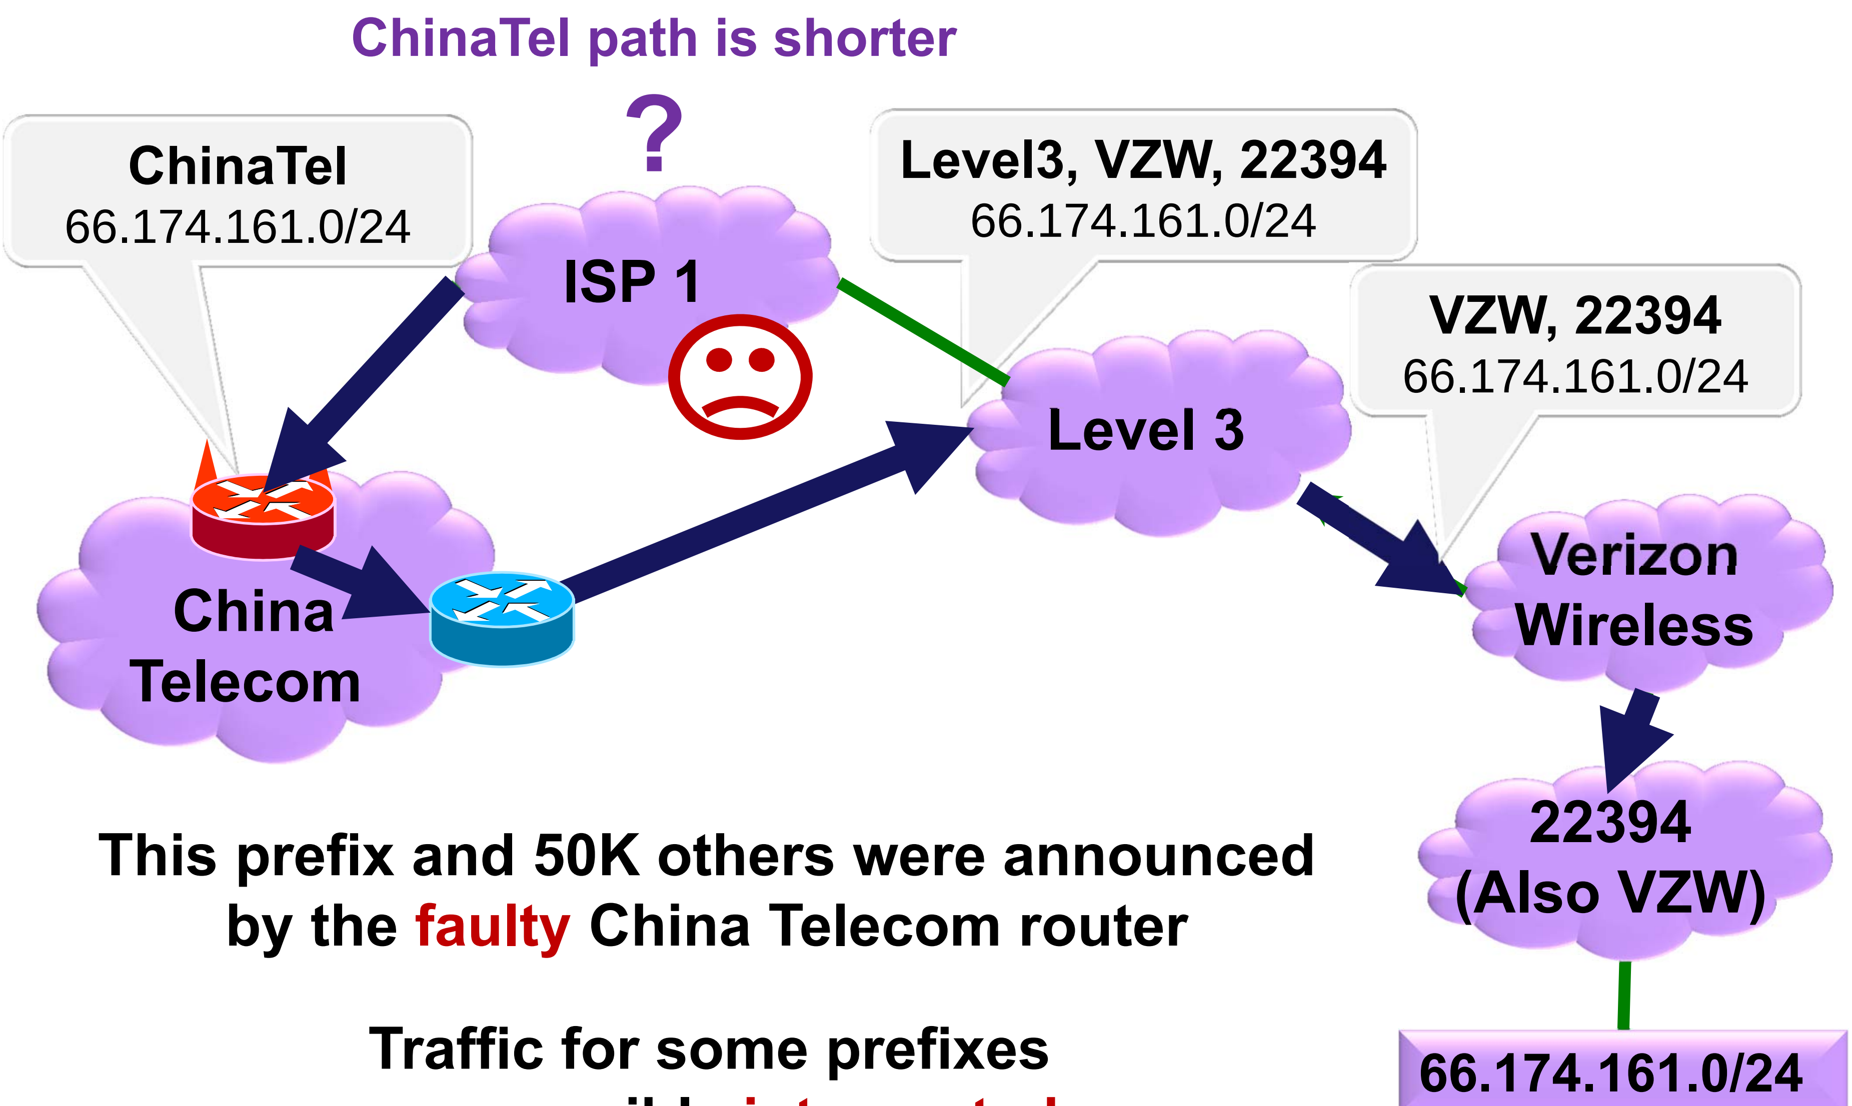

This prefix and 50K others were announced by the faulty China Telecom router

Traffic for some prefixes was possibly intercepted

S*BGP (*e.g.*, BGPsec) can stop this attack

BGPsec:

RPKI +

Cannot announce a path that was not announced to you.

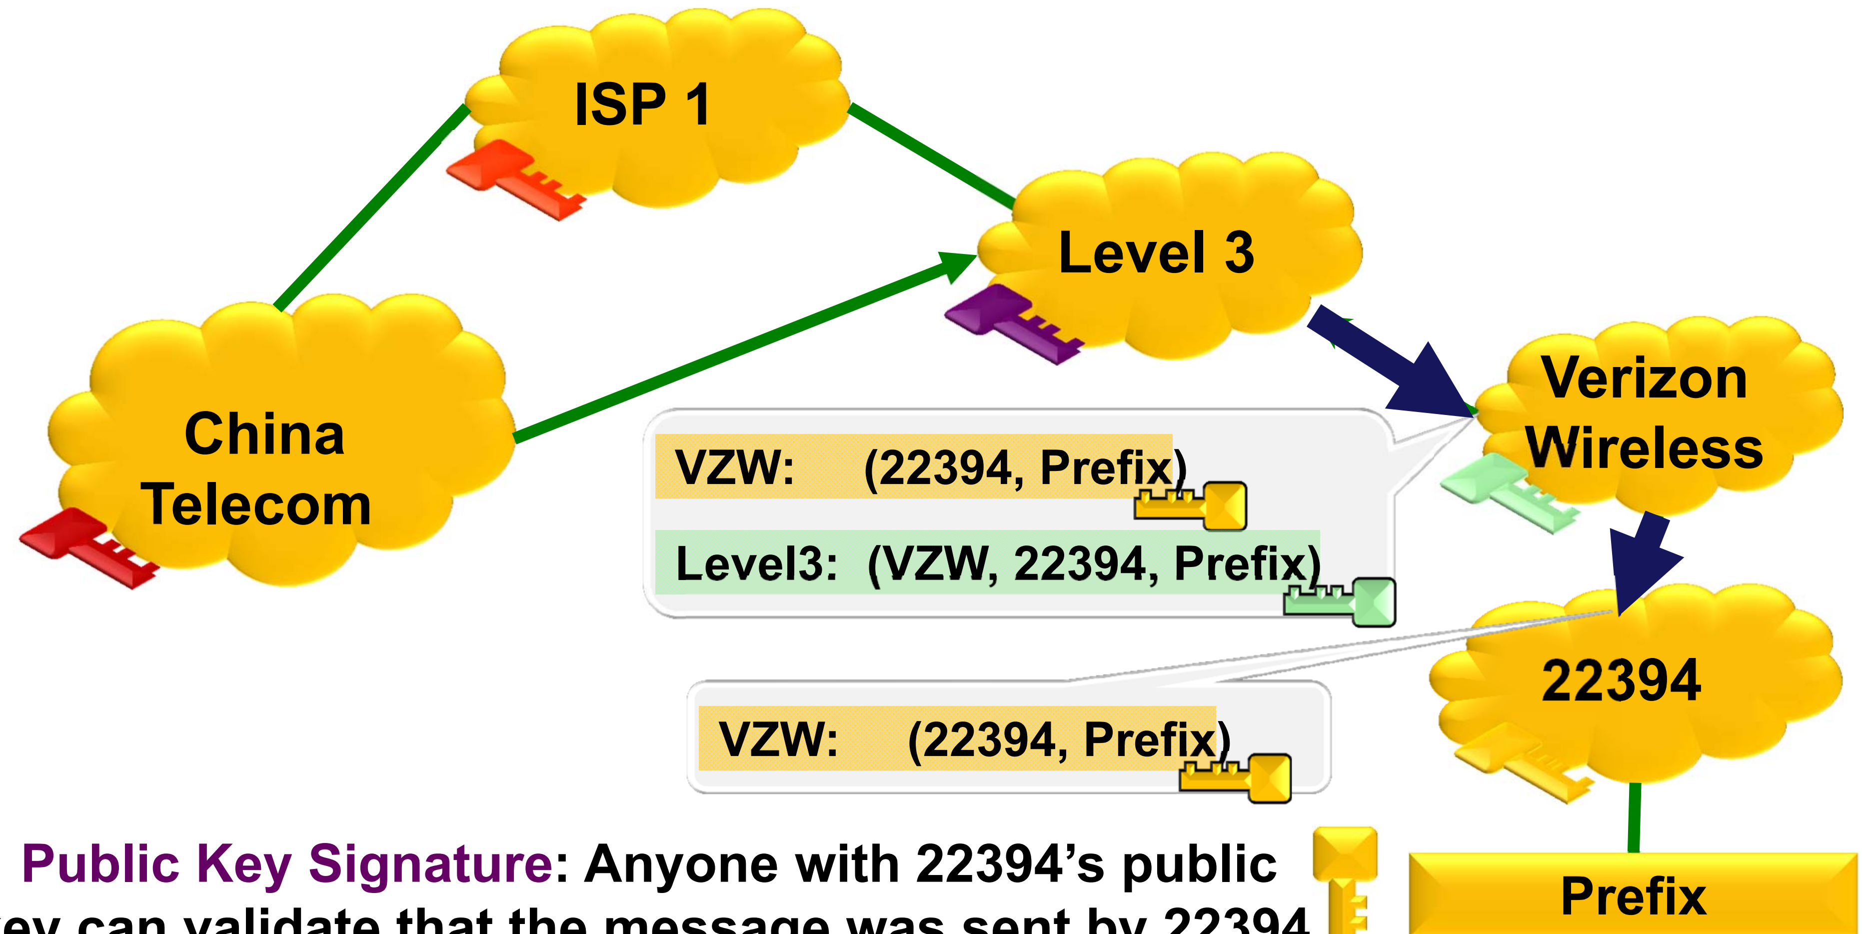

Public Key Signature: Anyone with 22394's public key can validate that the message was sent by 22394.

S*BGP (*e.g.*, BGPsec) can stop this attack

BGPsec:

Cannot announce a path that w

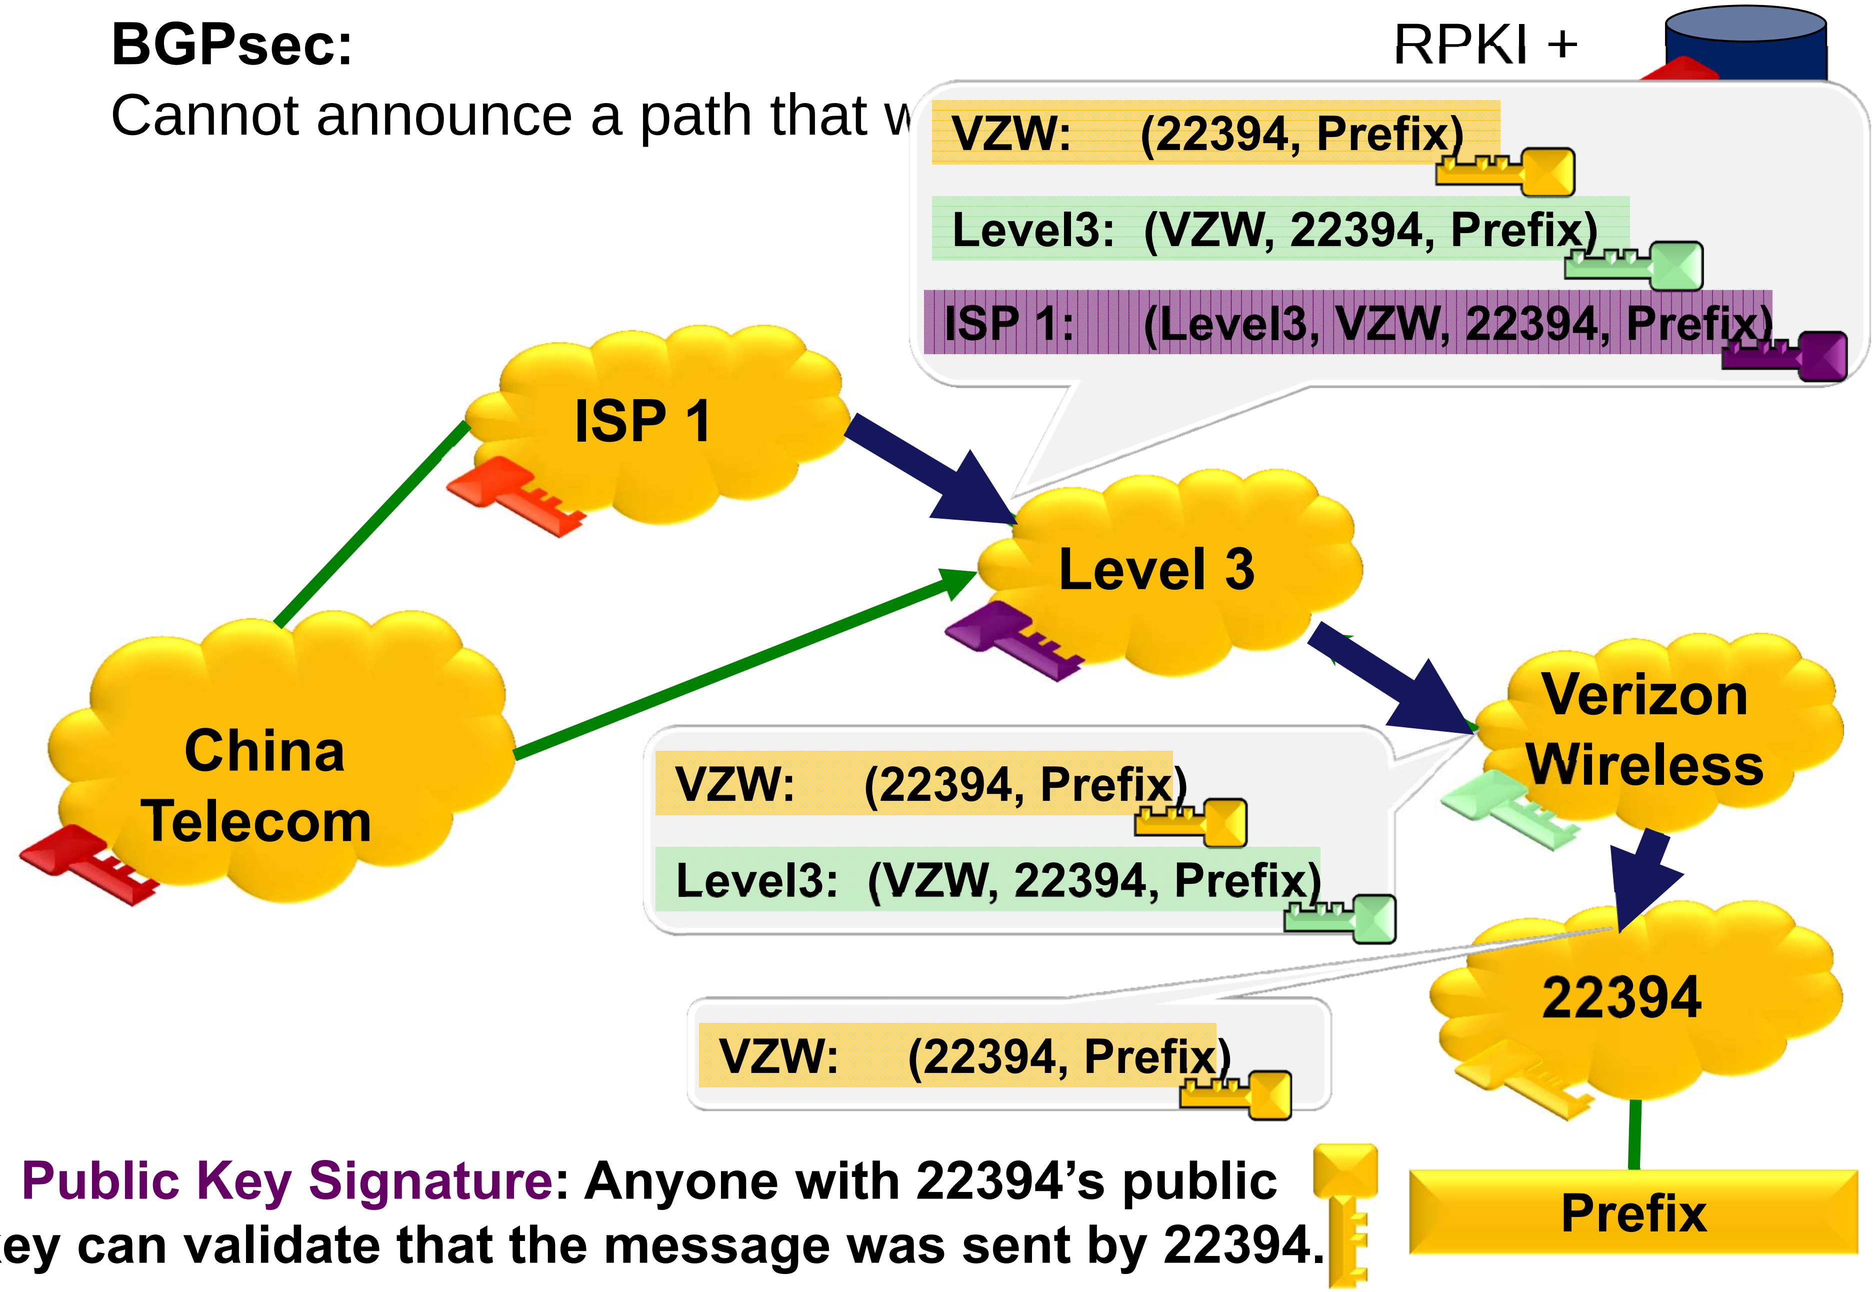

Public Key Signature: Anyone with 22394's public key can validate that the message was sent by 22394.

But what happens in partial deployment?

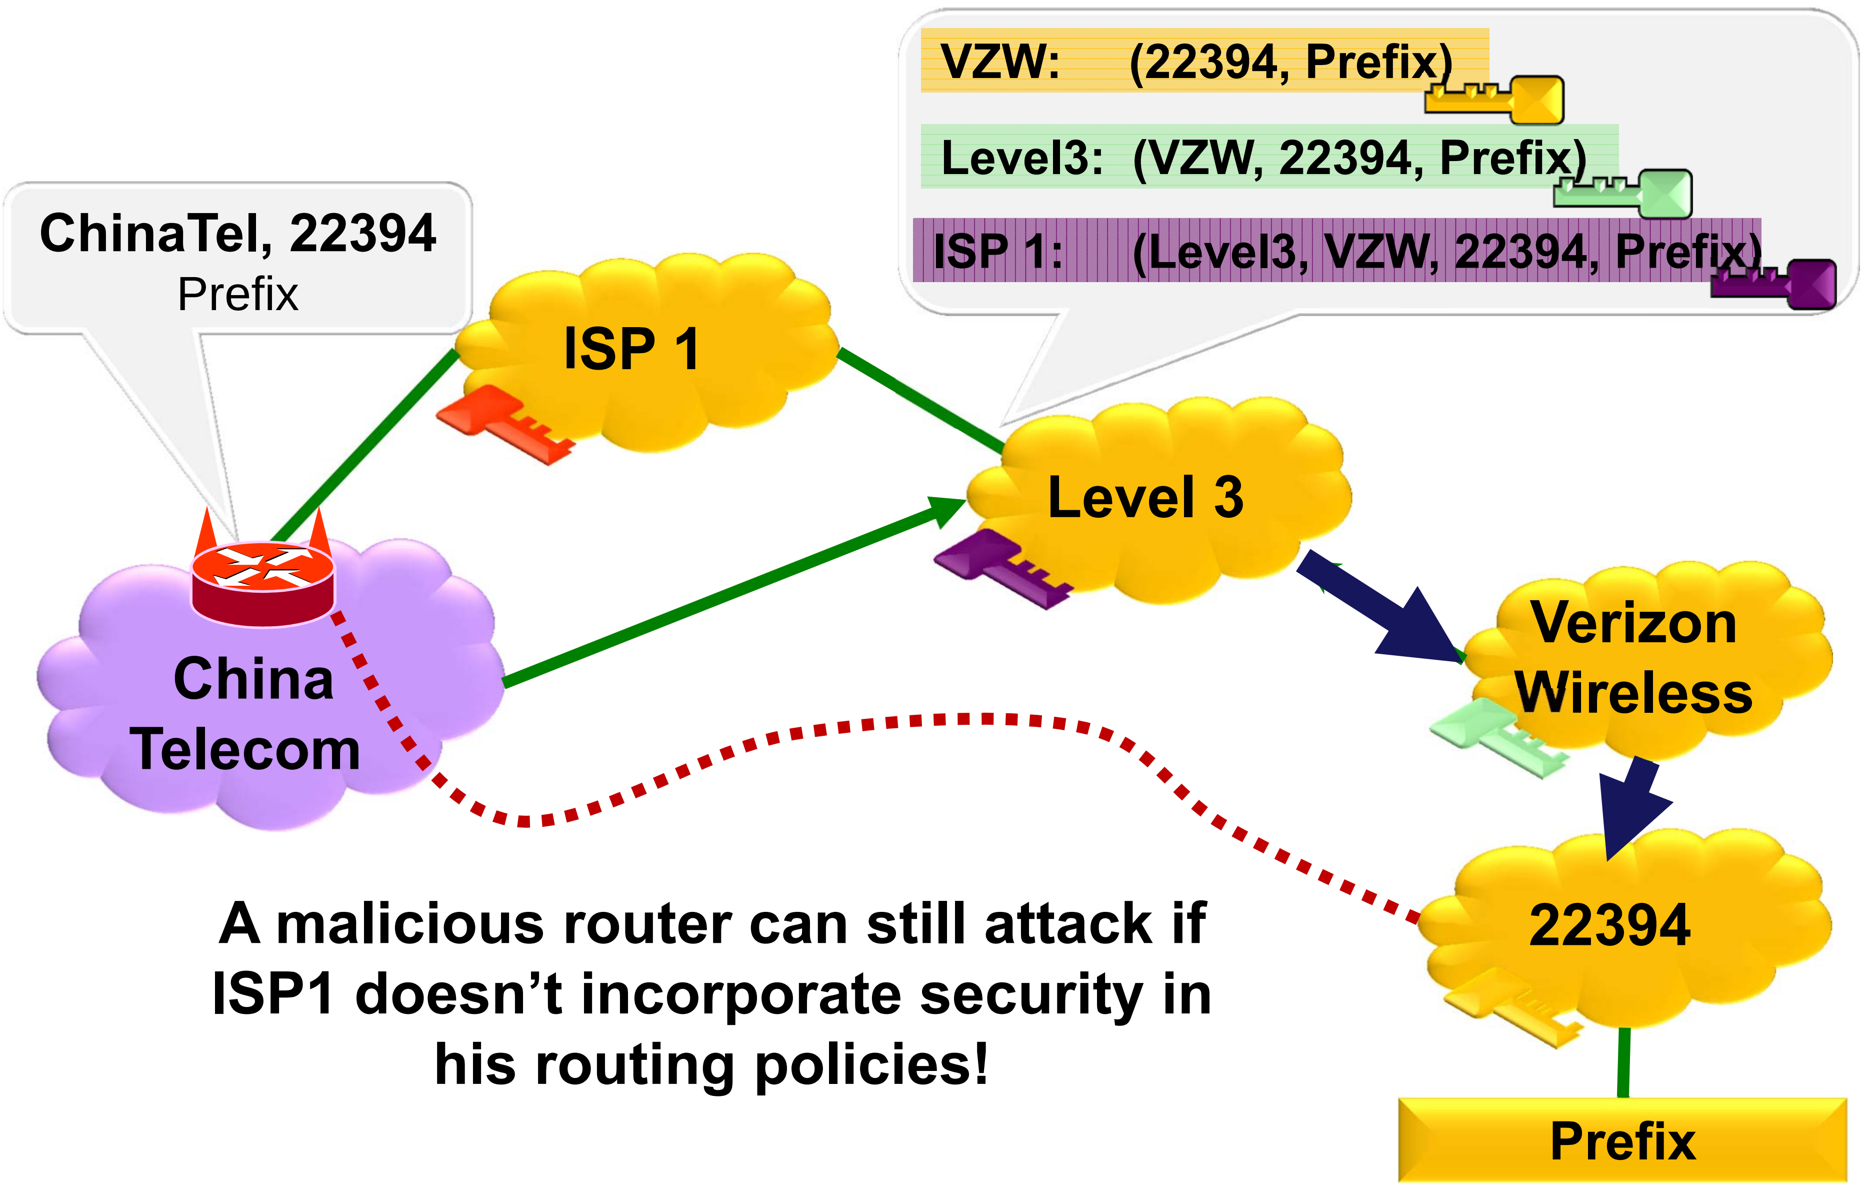

A malicious router can still attack if ISP1 doesn't incorporate security in his routing policies!

But what happens in partial deployment?

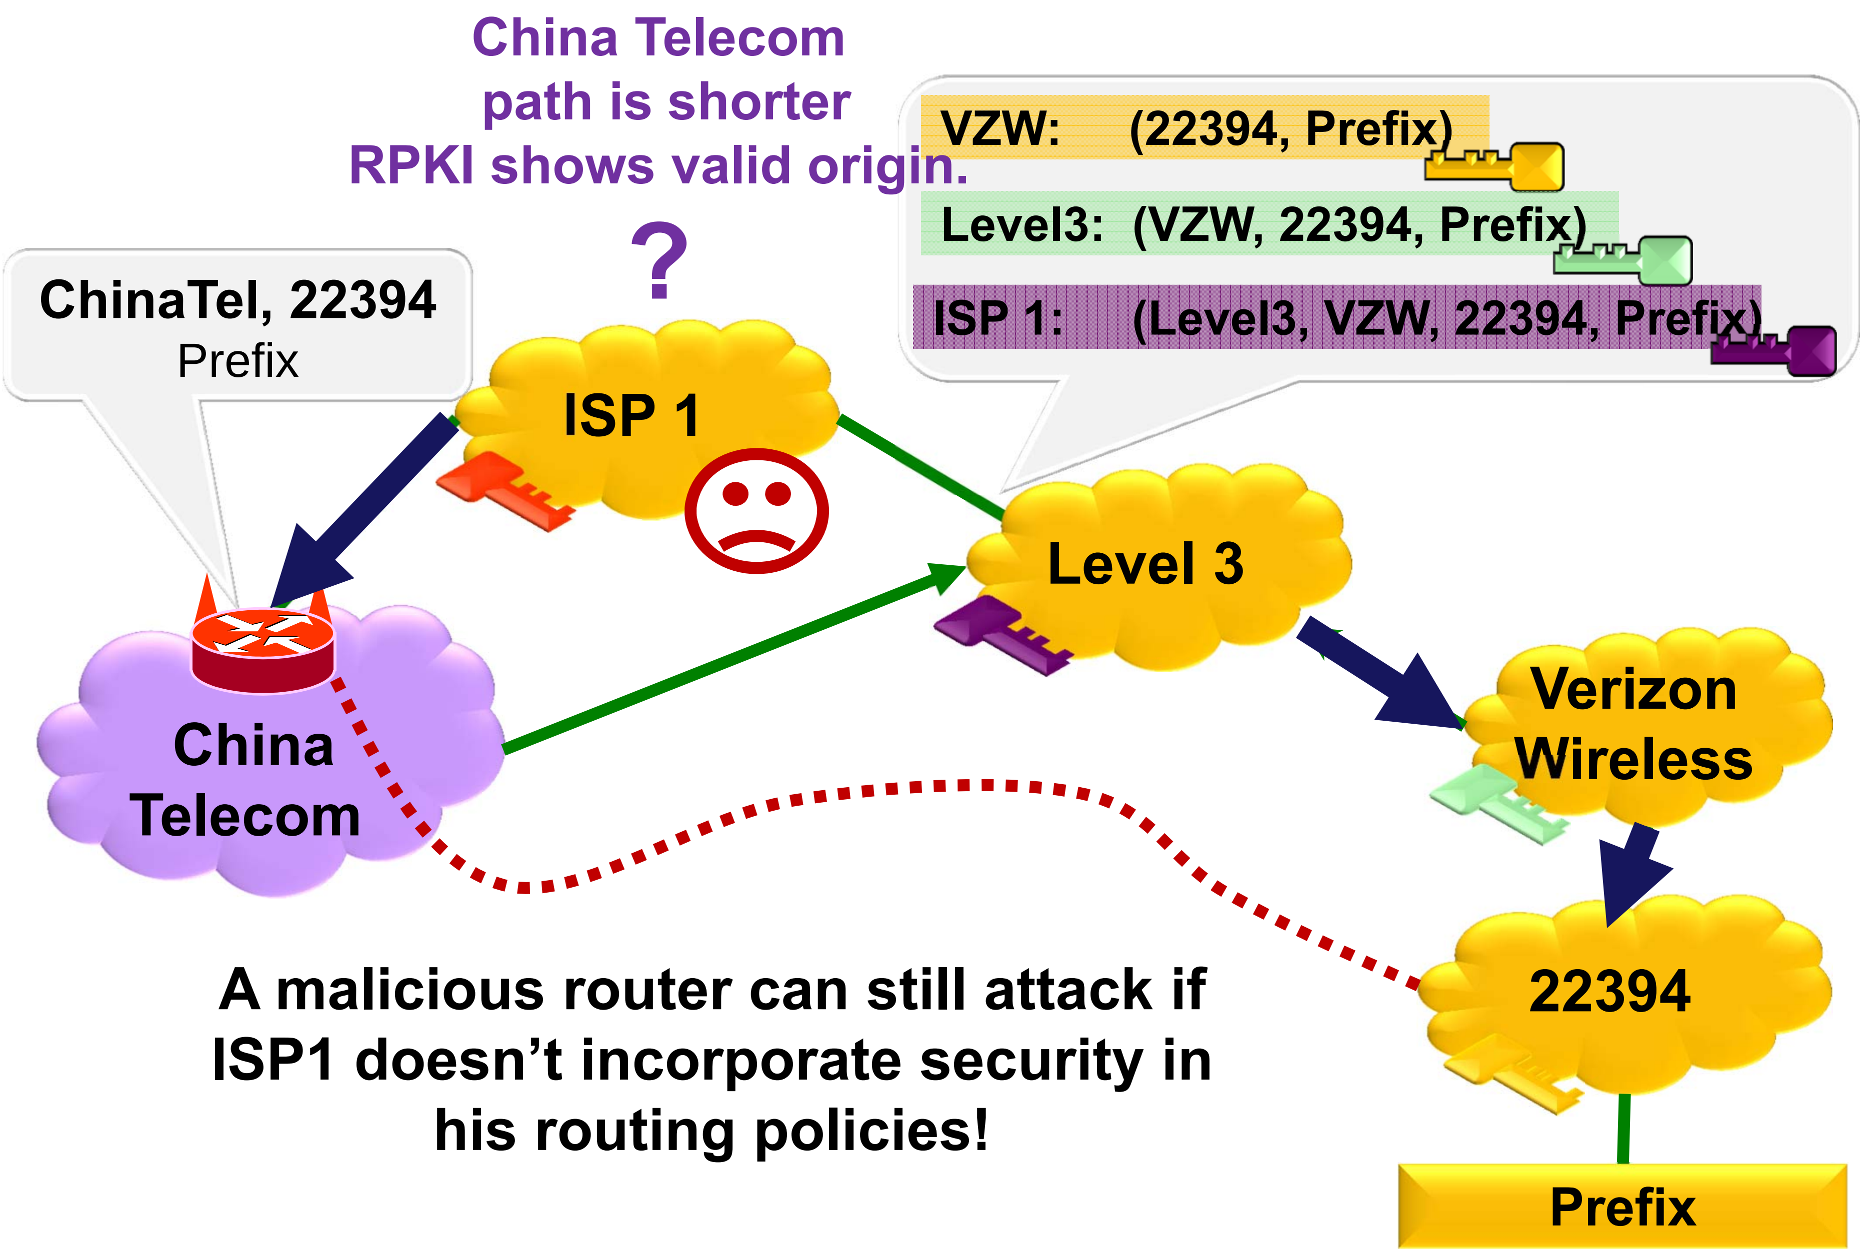

A malicious router can still attack if ISP1 doesn't incorporate security in his routing policies!

Bottom Line:

It’s not enough just to sign & validate!

If we want security, S*BGP must also impact BGP routing policy.

Key idea: S*BGP impacts routing & revenue (1)

Ideally (security geek):

Routing Policy:

1. **Prefer secure routes**
2. Local Pref
3. AS path
4. ….
5. Arbitrary tiebreaks

I know you don't like changing routing policies this much, so instead we assumed:

Routing Policy:

1. Local Pref
2. AS path
3. ….
4. **Prefer secure routes**
5. Arbitrary tiebreaks

***i.e.,* Secure ISPs *at least* break ties in favor of secure routes.**

Talk Organization

Background:

- BGP, attacks and defenses like **RPKI** & **BGPsec**

A Strategy for S*BGP deployment

Evaluating our strategy

1. Model
2. Simulation results on **[UCLA Cyclops]** AS graph data

Conclusions and recommendations

Let’s switch gears….

Instead of security,
let’s start thinking about economics.

AS-level Business Relationships

A simple model of AS-level business relationships.

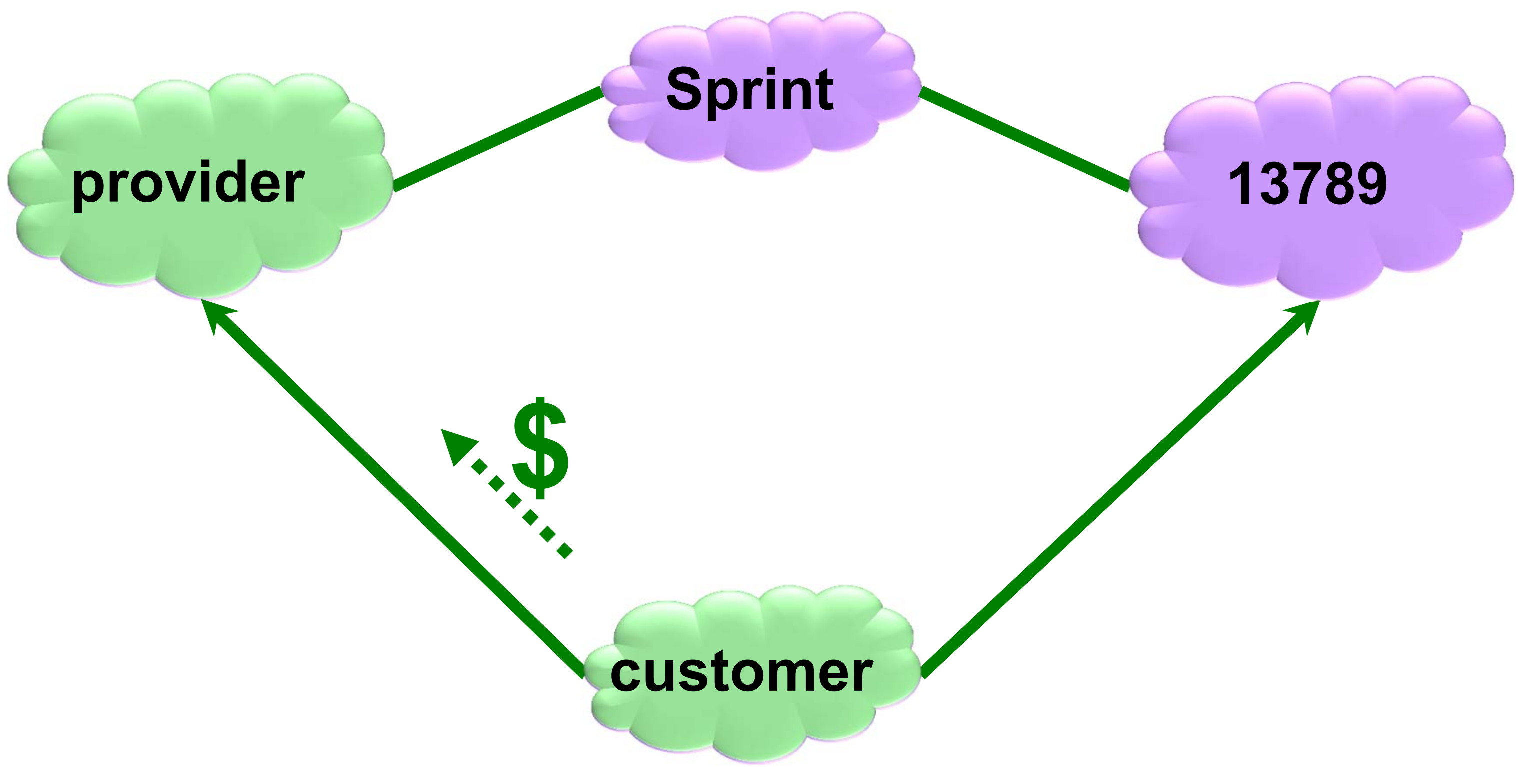

AS-level Business Relationships

A simple model of AS-level business relationships.

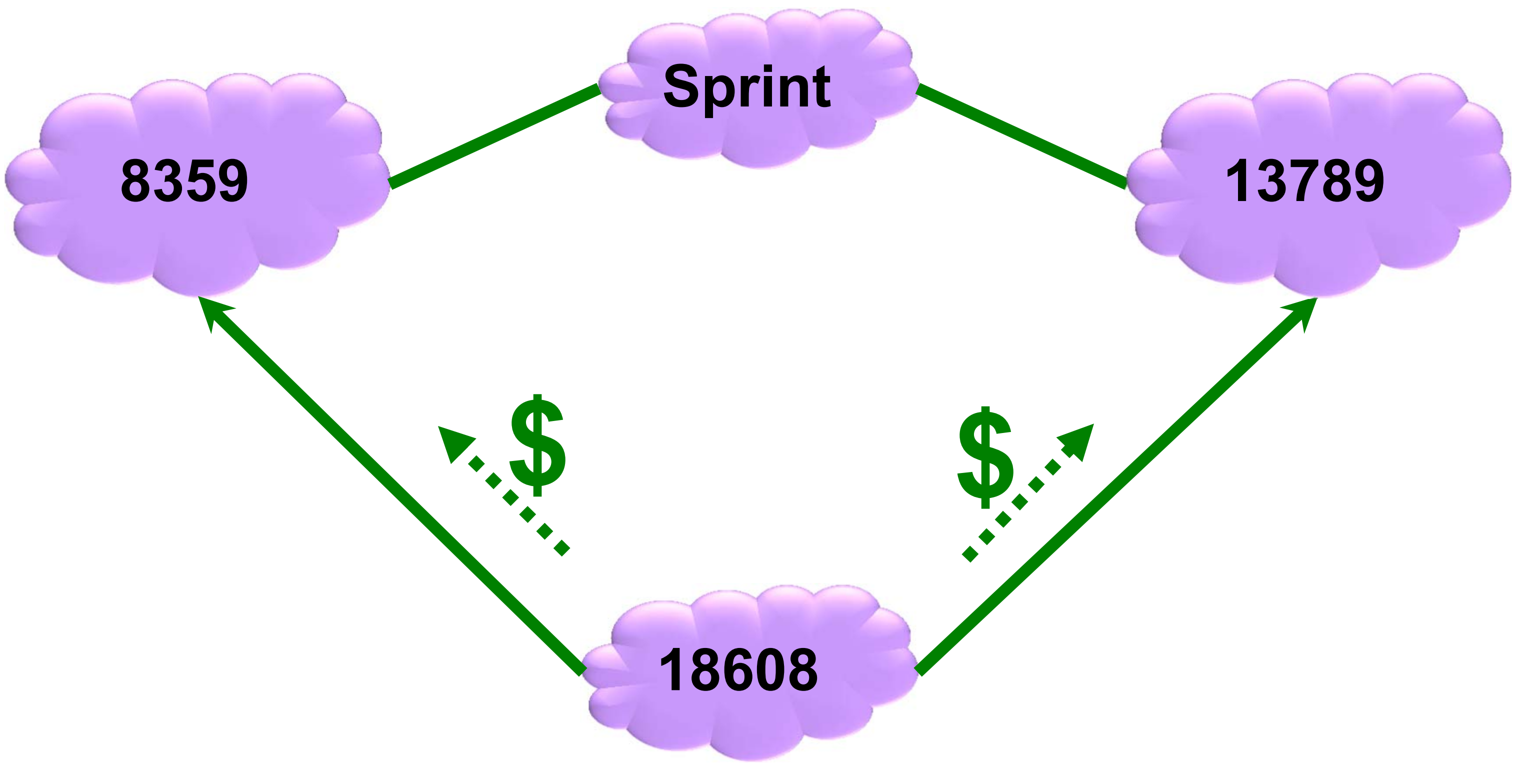

AS-level Business Relationships

A simple model of AS-level business relationships.

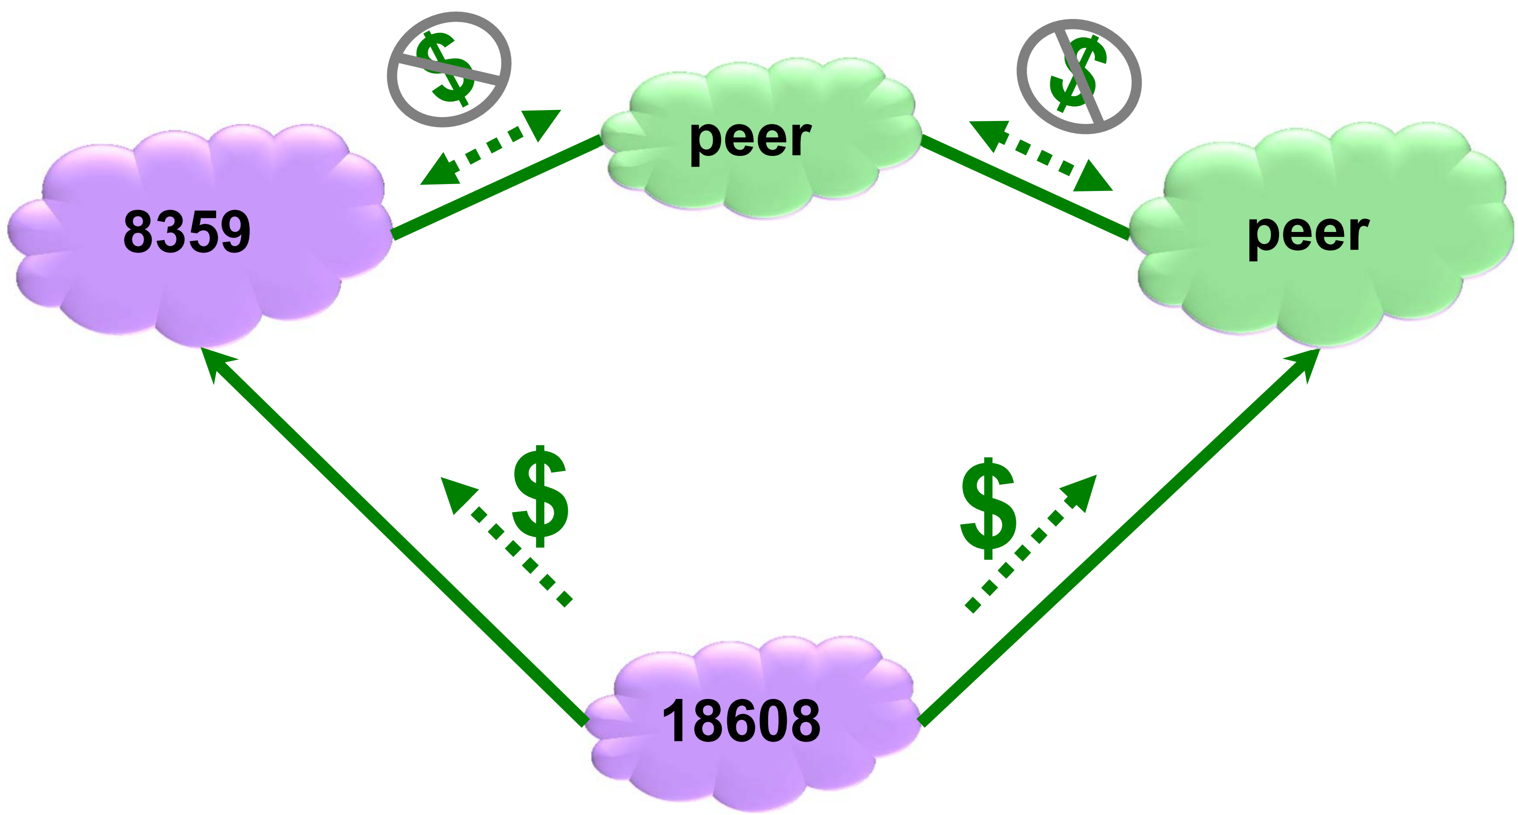

Stubs vs ISPs: Stubs are 85% of the Internet's ASes!

A stub is an AS with no customers.

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

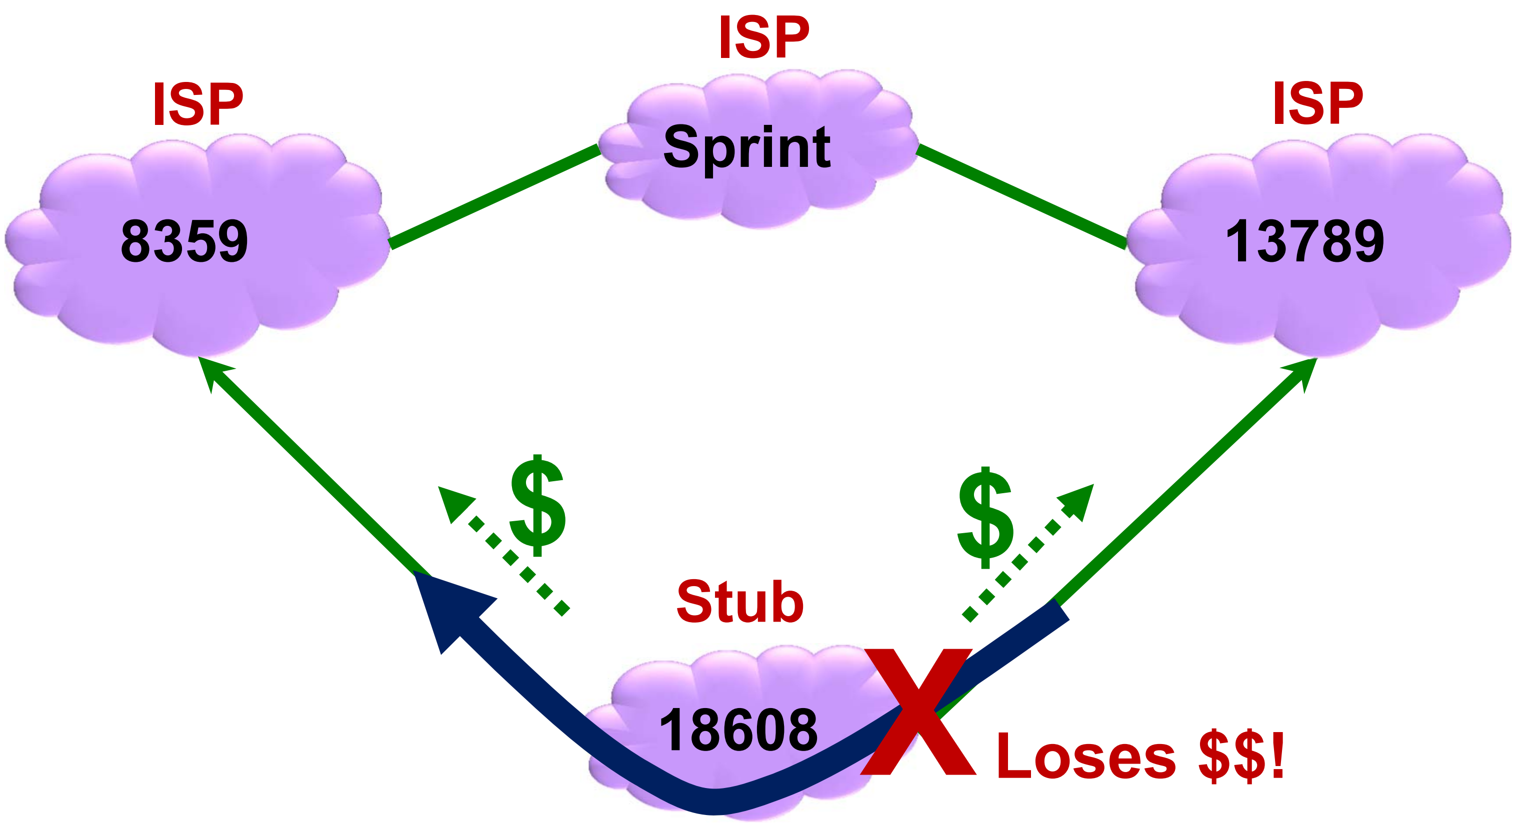

85% of ASes are stubs! We call the rest (15%) ISPs.

How to drive S*BGP deployment:

An ISP attracts more customer traffic

=

It earns more revenue

Key idea: S*BGP impacts routing & thus revenue (1)

Assume: Secure ISPs *at least* break ties in favor of secure paths

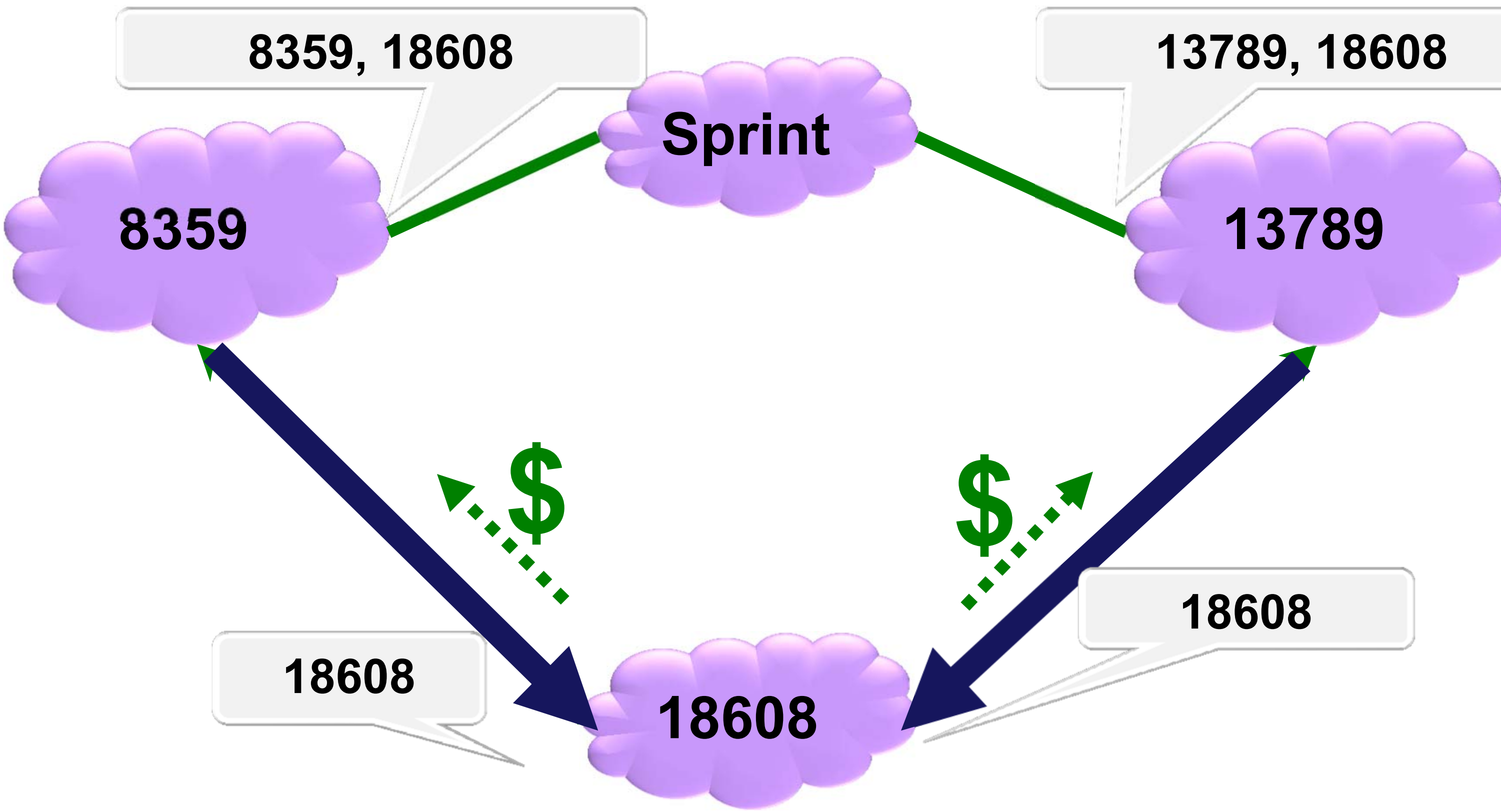

ISPs can use S*BGP to attract customer traffic & thus money

Key idea: S*BGP impacts routing & thus revenue (1)

Assume: Secure ISPs *at least* **break ties** in favor of secure paths

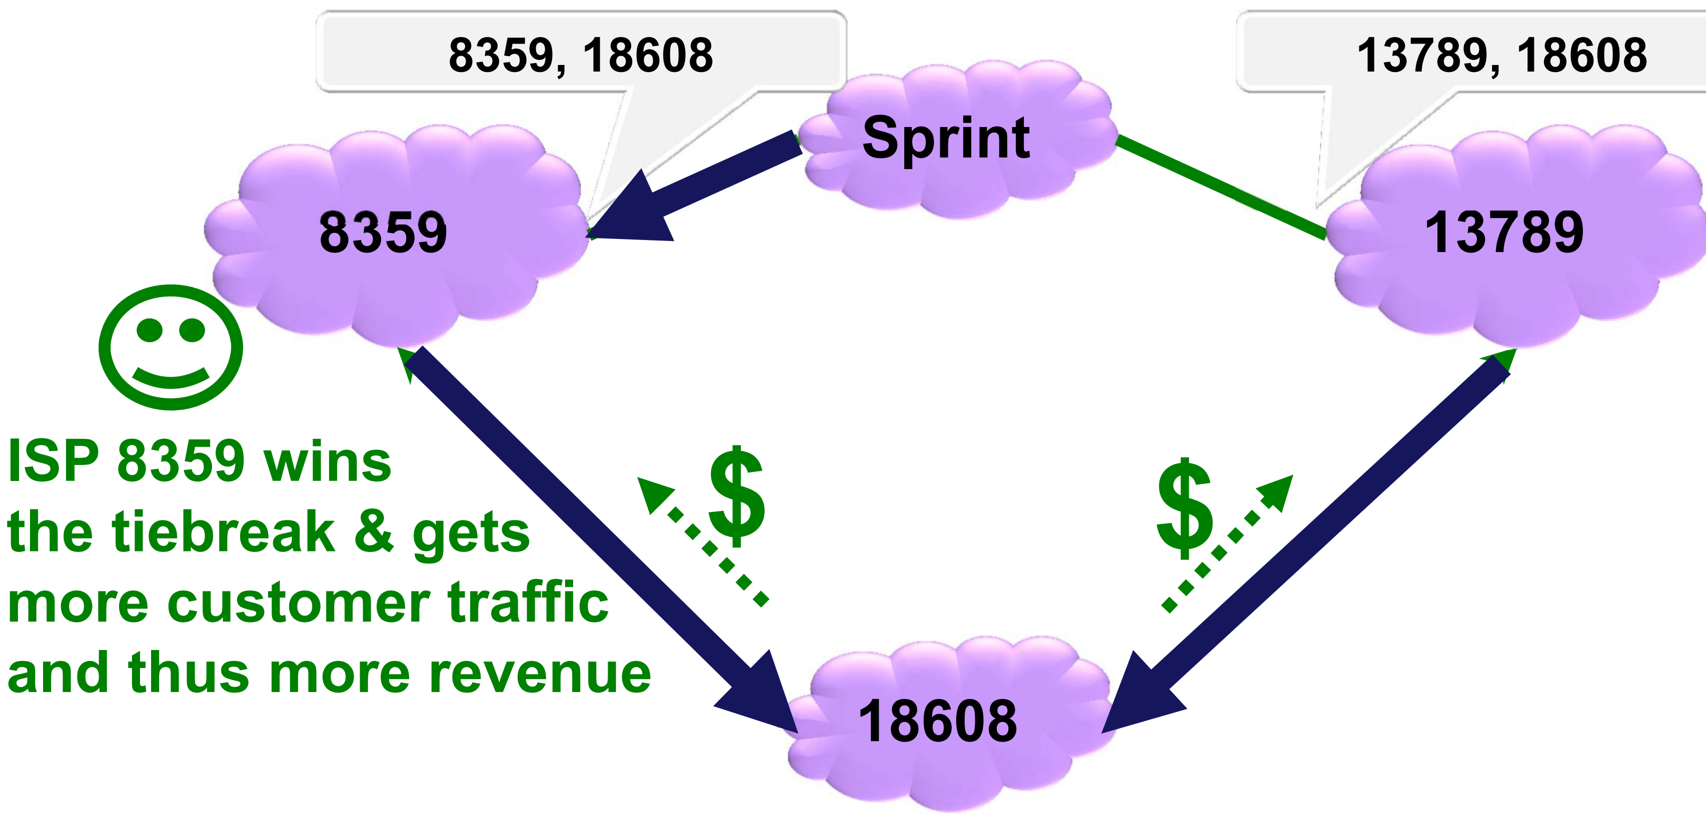

ISPs can use S*BGP to attract customer traffic & thus money

Key idea: S*BGP impacts routing & thus revenue (2)

Assume: Secure ISPs *at least* break ties in favor of secure paths

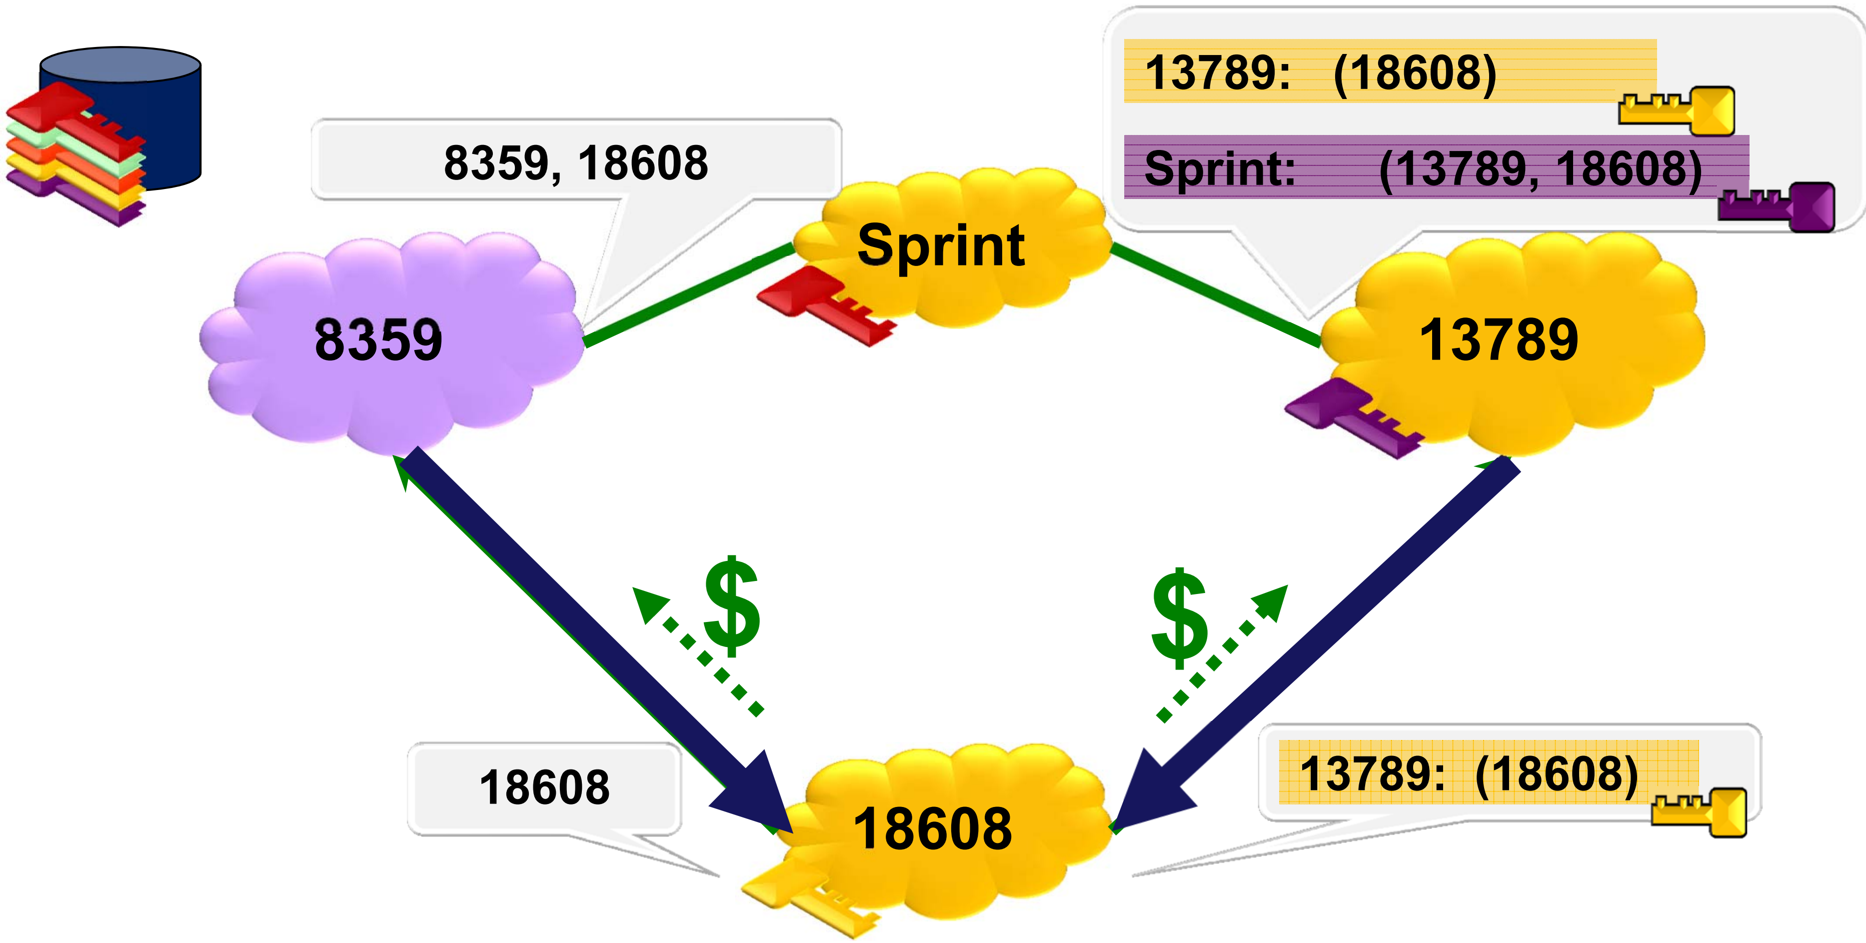

ISPs can use S*BGP to attract customer traffic & thus money

Key idea: S*BGP impacts routing & thus revenue (2)

Assume: Secure ISPs *at least* break ties in favor of secure paths

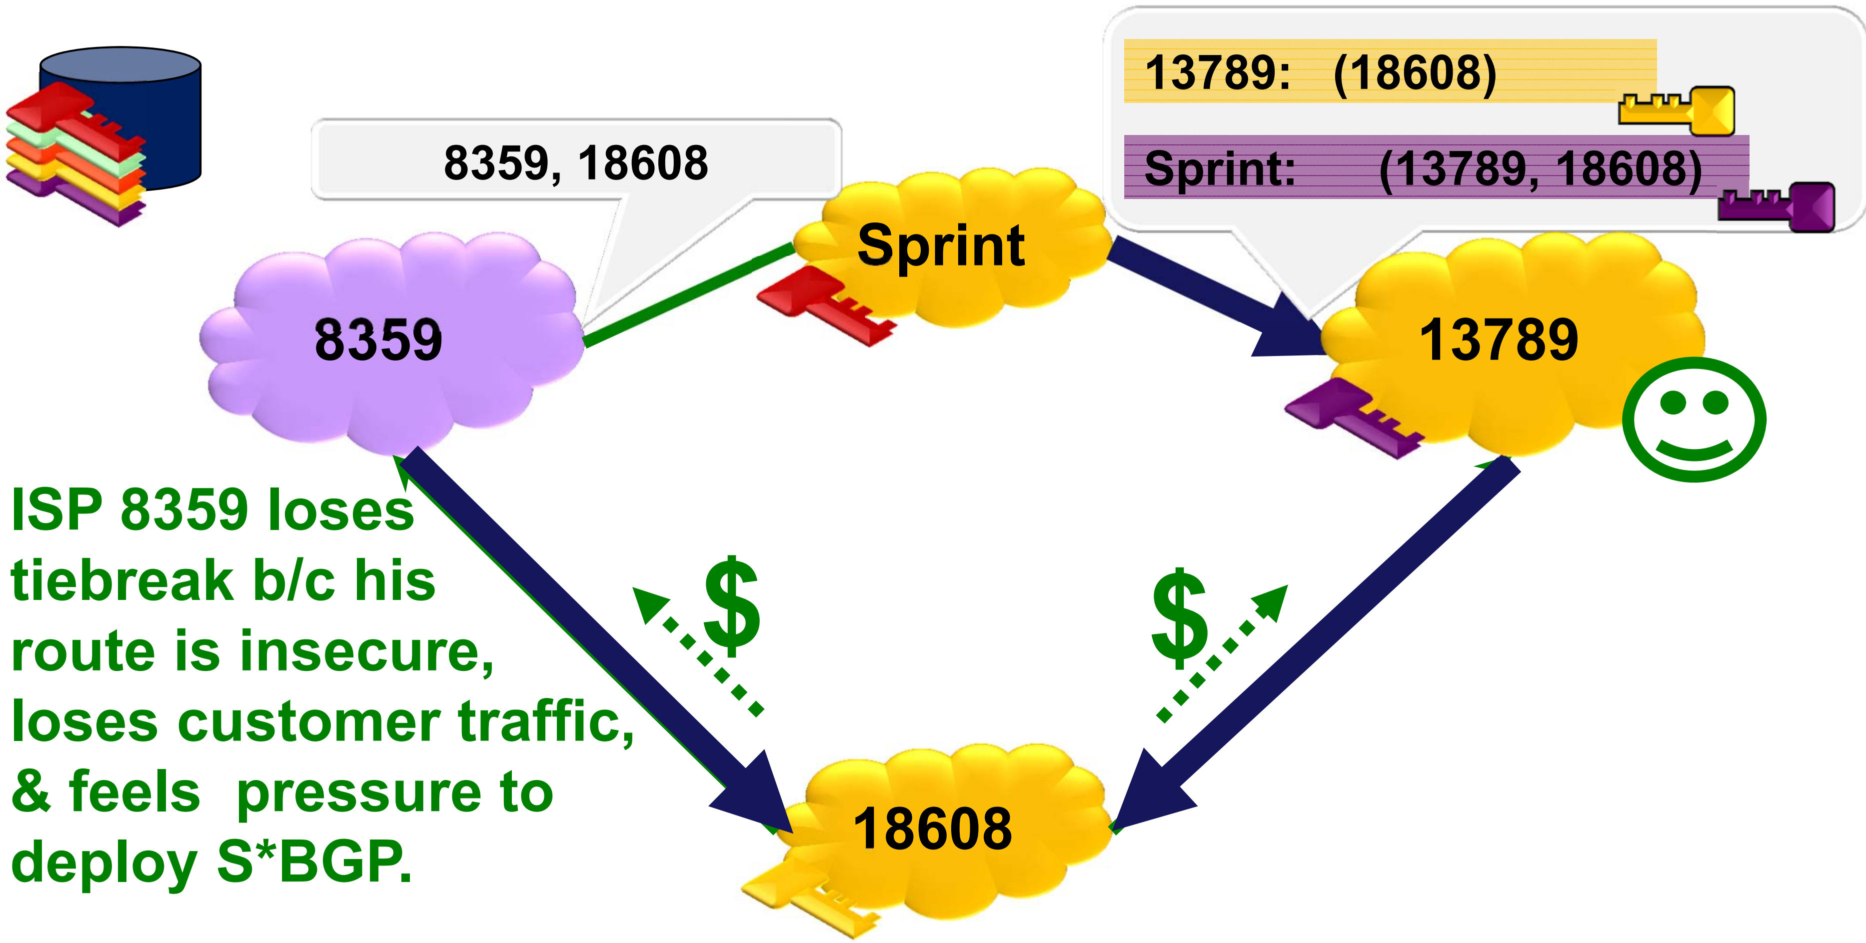

ISP 8359 loses tiebreak b/c his route is insecure, loses customer traffic, & feels pressure to deploy S*BGP.

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

Our Main Result: A Strategy for Deploying S*BGP

1. **Secure ISPs *at least* break ties in favor of secure paths**

2. **A few early adopters initially deploy S*BGP (gov't incentives, regulations, security concerns, etc.)**

   (A least 5 of the biggest Tier 1s)

3. **ISPs deploy simplex S*BGP in their stub customers**

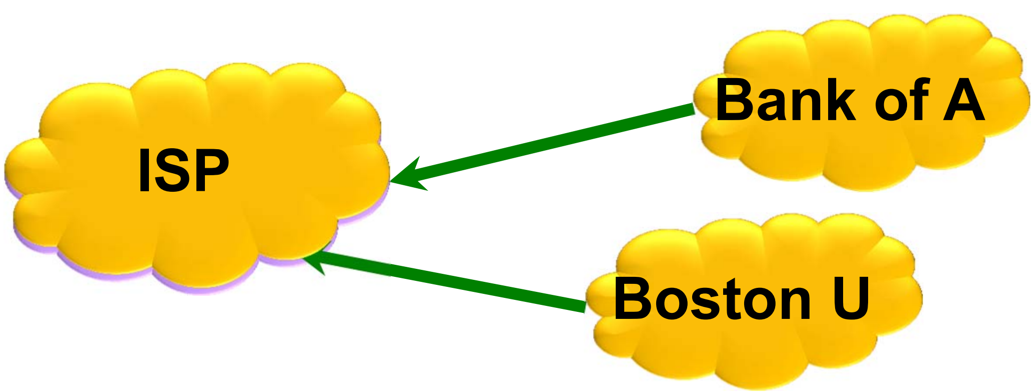

Stub with Simplex S*BGP:

- Need only sign; trusts provider to validate.
- Minor security impact
- No hardware upgrade!

(Gov'ts should subsidize ISPs that do this.)

Crucial, since 85% of ASes are stubs!

Talk Organization

Background:

- BGP, attacks and defenses like **RPKI** & **BGPsec**

A Strategy for S*BGP deployment

Evaluating our strategy

1. Model
2. Simulation results on **[UCLA Cyclops]** AS graph data

Conclusions and recommendations

A model of the S*BGP deployment process

- **To start the process:**
  - **Early adopter ASes** become secure
  - Their **stub** customers become secure (e.g. simplex S*BGP)

- **Each round:**
  - Compute **customer traffic volume** for **every insecure** ISP n
  - If ISP n 's customer traffic 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 "customer traffic volume"?

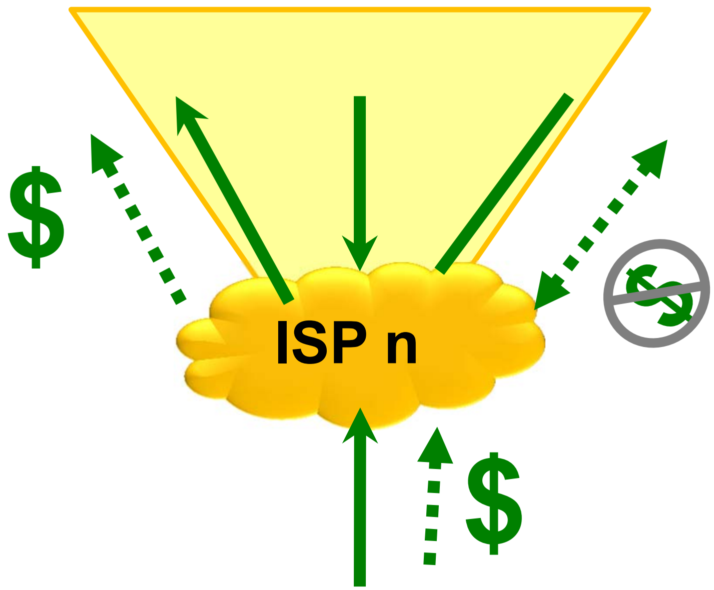

Number of source ASes
routing through ISP n
to all customer destinations.

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

Talk Organization

Background:

- BGP, attacks and defenses like **RPKI** & **BGPsec**

A Strategy for S*BGP deployment

Evaluating our strategy

1. Model
2. Simulation results on **[UCLA Cyclops]** data

Conclusions and recommendations

Simulation: Deployment Case Study

Ten early adopters:

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

- The five content providers source **10%** of all Internet traffic
- All nodes have the same threshold **θ = 5%.**

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

Bottom Line:

It works.

Let’s see why…
with an excerpt from our simulations.

Simulation: Market pressure drives deployment (1)

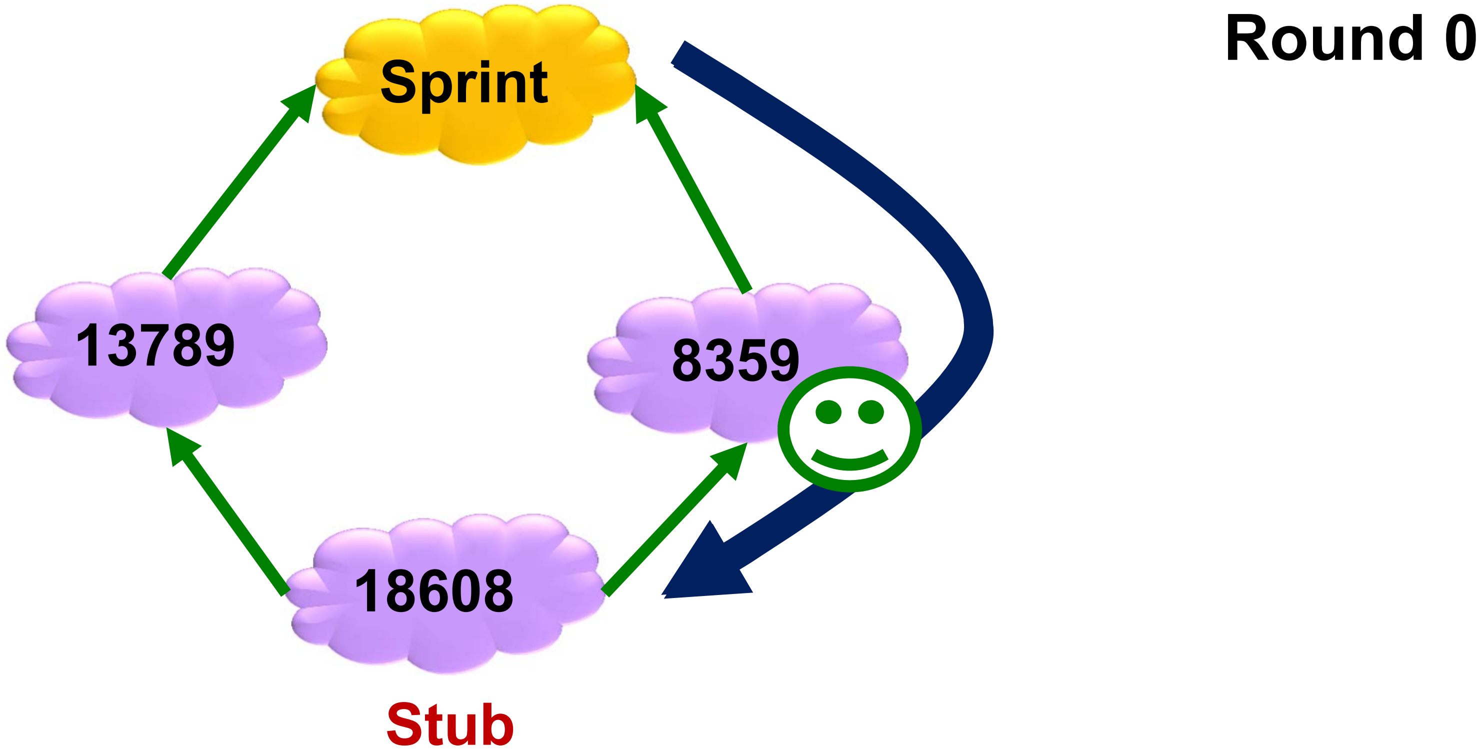

Notice that Sprint is offered two equally good (customer, 2 hop) paths to stub AS18608. The tiebreak algorithm prefers AS 8359.

AS8359 is happy because he gets revenue from traffic from Sprint to AS18608.

Simulation: Market pressure drives deployment (1)

Round 1

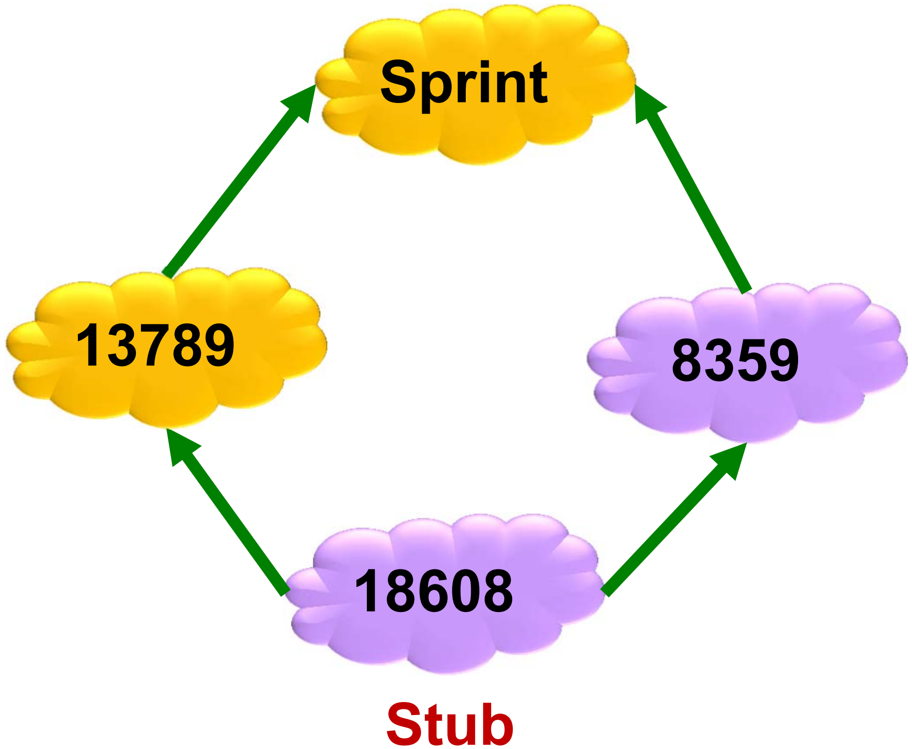

Simulation: Market pressure drives deployment (1)

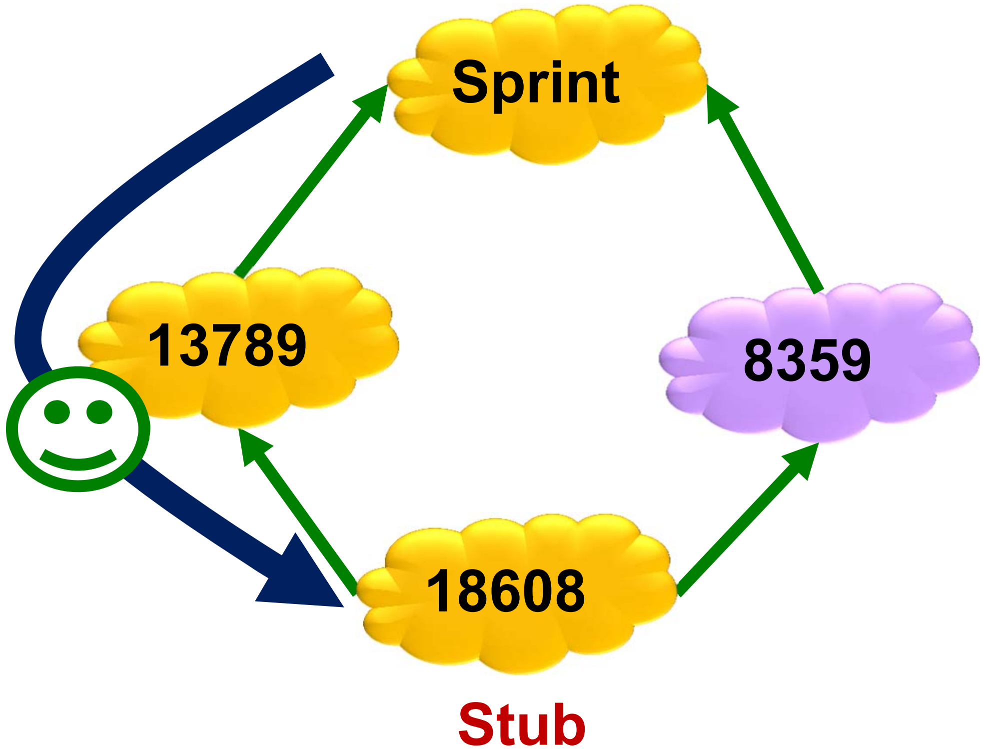

Now, AS 13789 deploys S*BGP in himself and his stub to draw traffic away from AS 8359.

(This is only a fraction of the traffic AS 13789 steals from competitors; for clarity we only show a small subgraph where he steals traffic here. Remember we compute traffic flow to ALL 36K ASes in the Internet, so AS13879 could have stolen traffic to many stubs.)

Simulation: Market pressure drives deployment (1)

Round 4

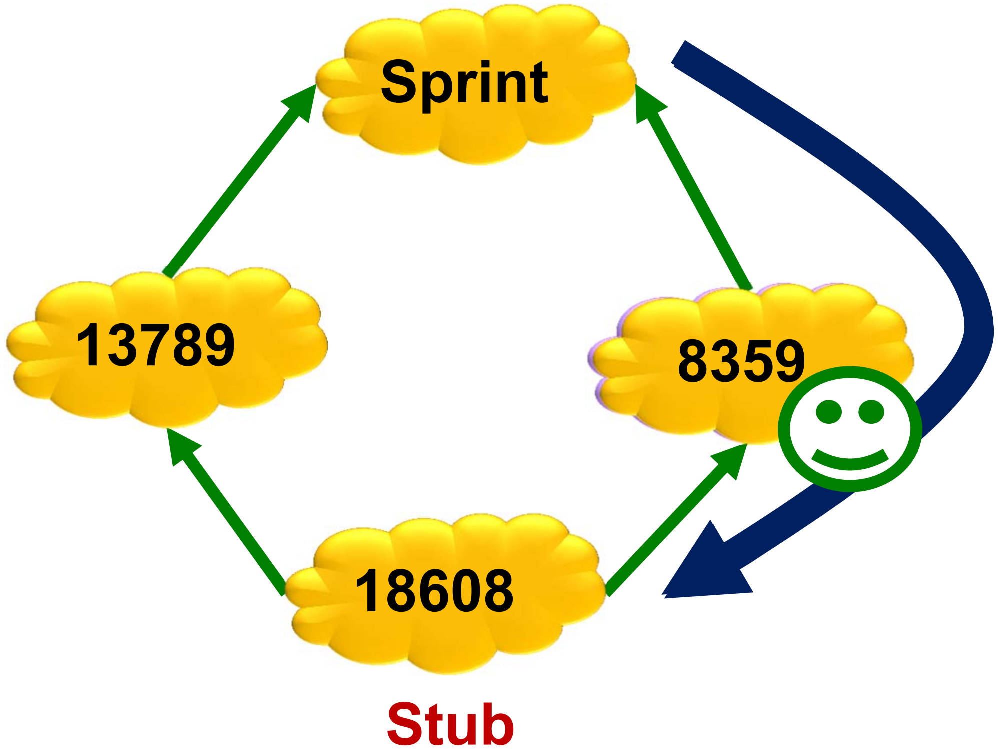

Now, AS 8359 deploys S*BGP to get back the traffic he lost to AS13789!

Simulation: Market pressure drives deployment (2)

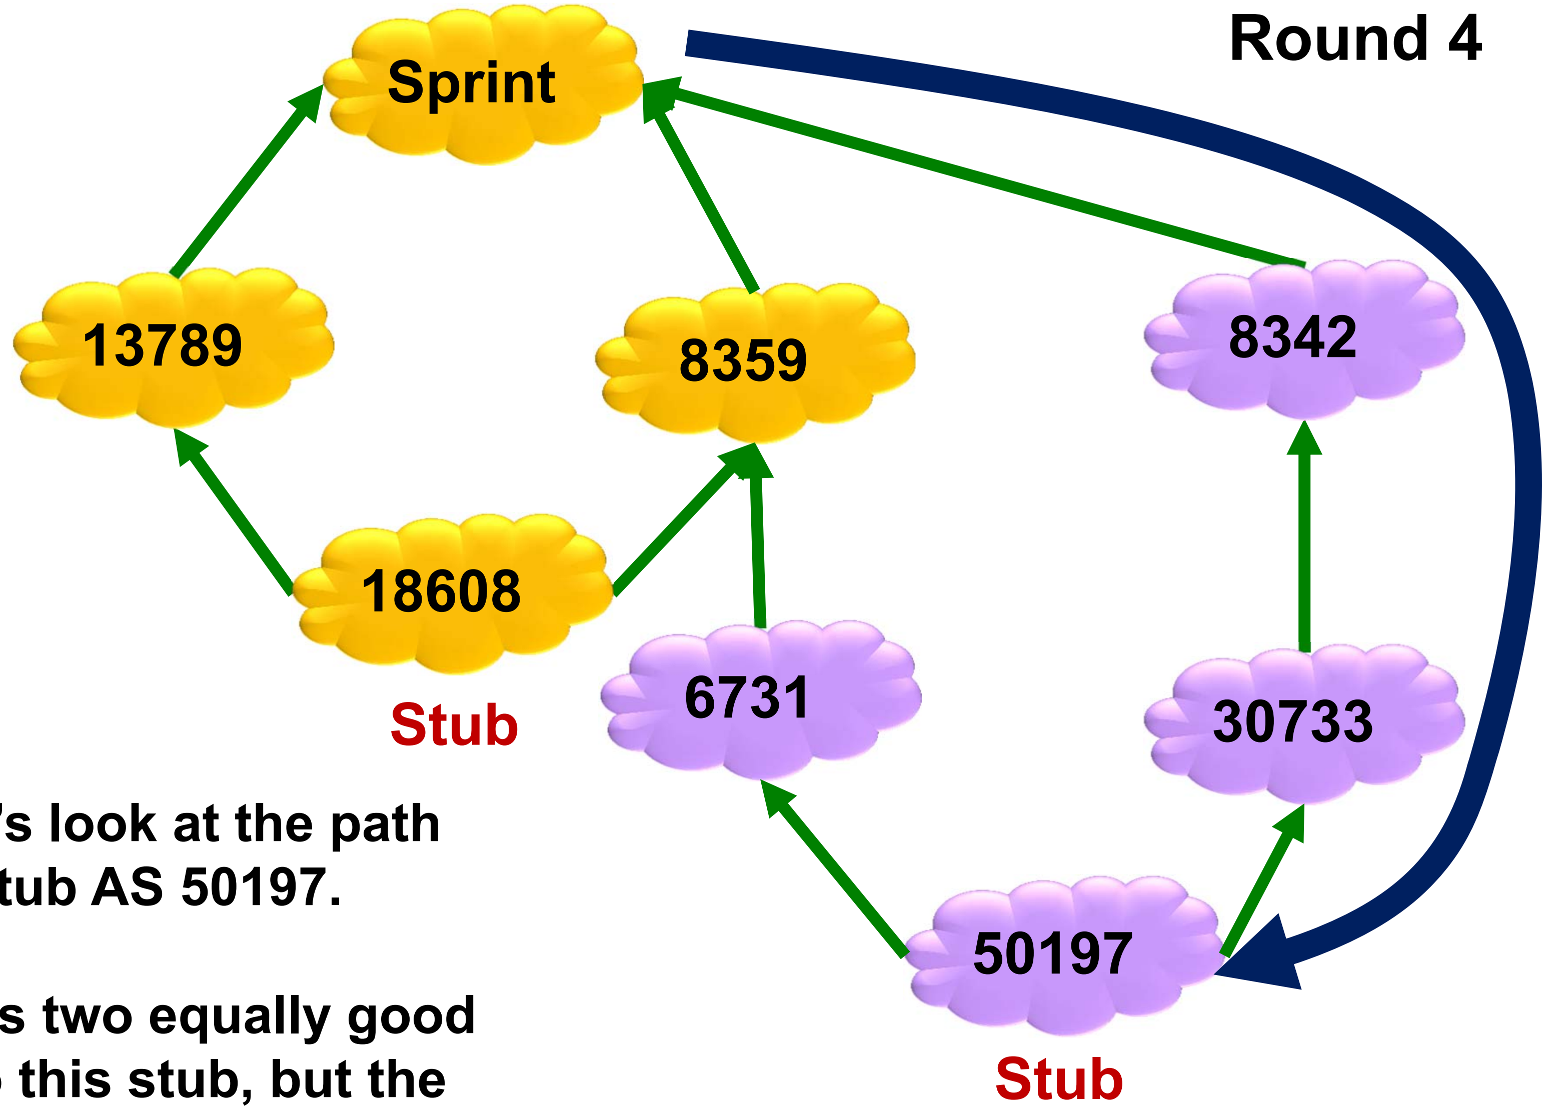

Now let's look at the path to stub AS 50197.

Sprint has two equally good paths to this stub, but the tiebreak algorithm prefers AS8342.

Simulation: Market pressure drives deployment (2)

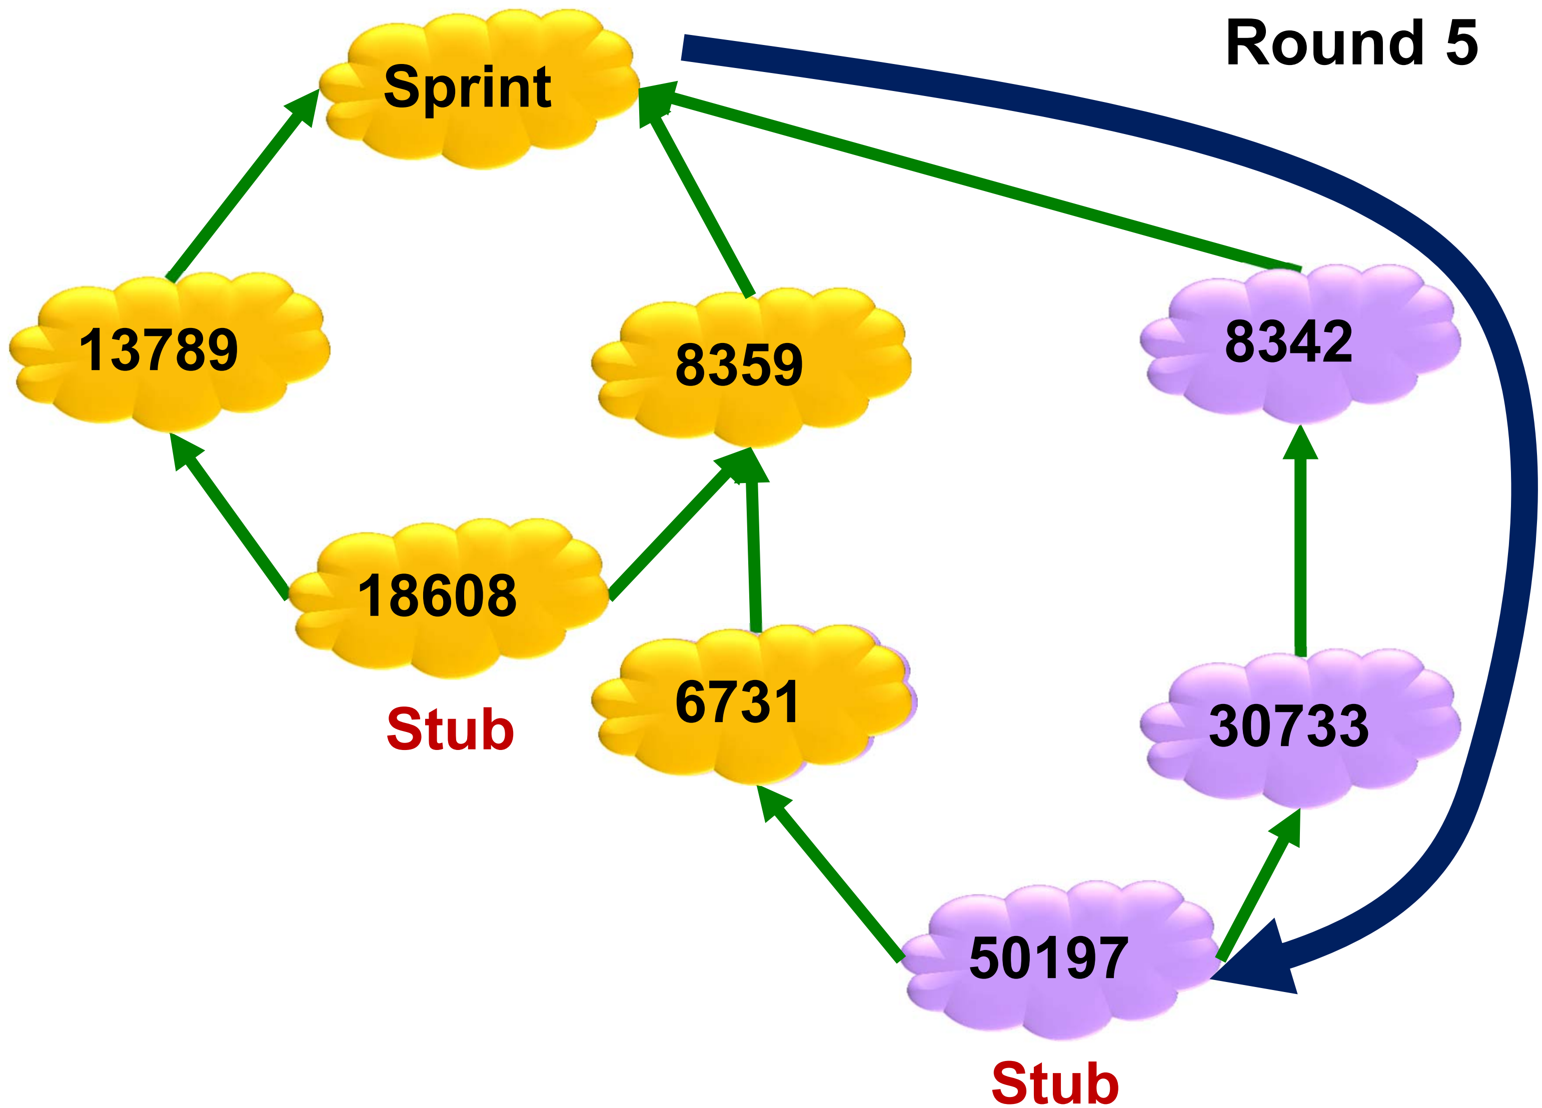

Simulation: Market pressure drives deployment (2)

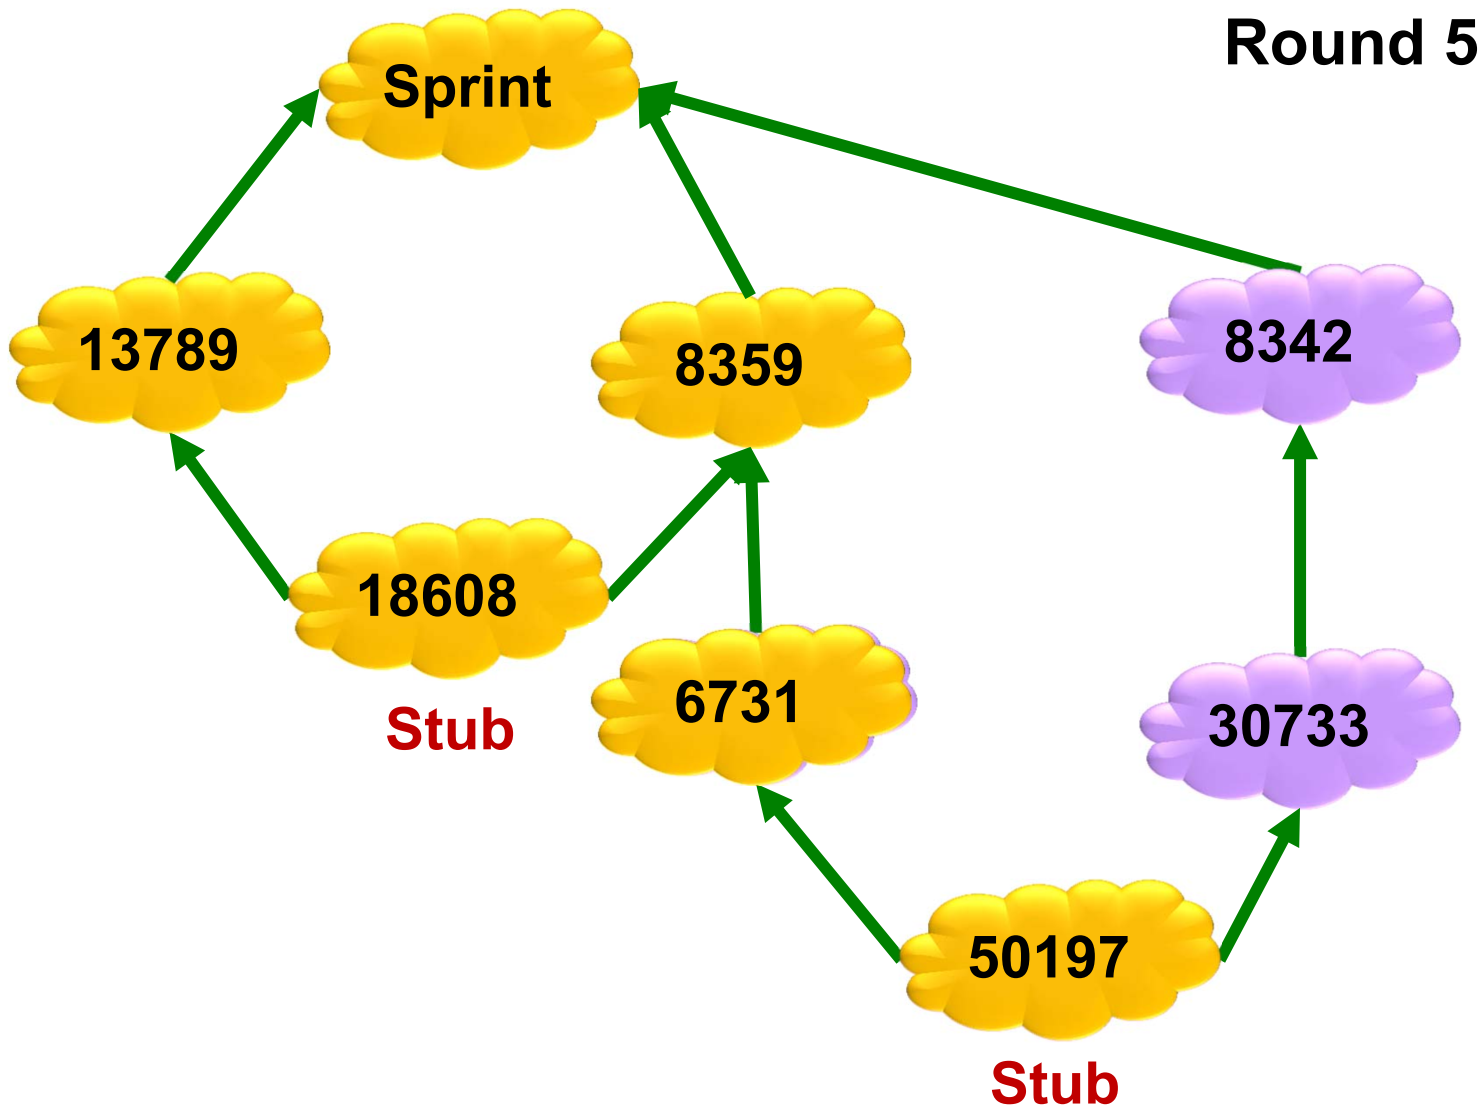

Simulation: Market pressure drives deployment (2)

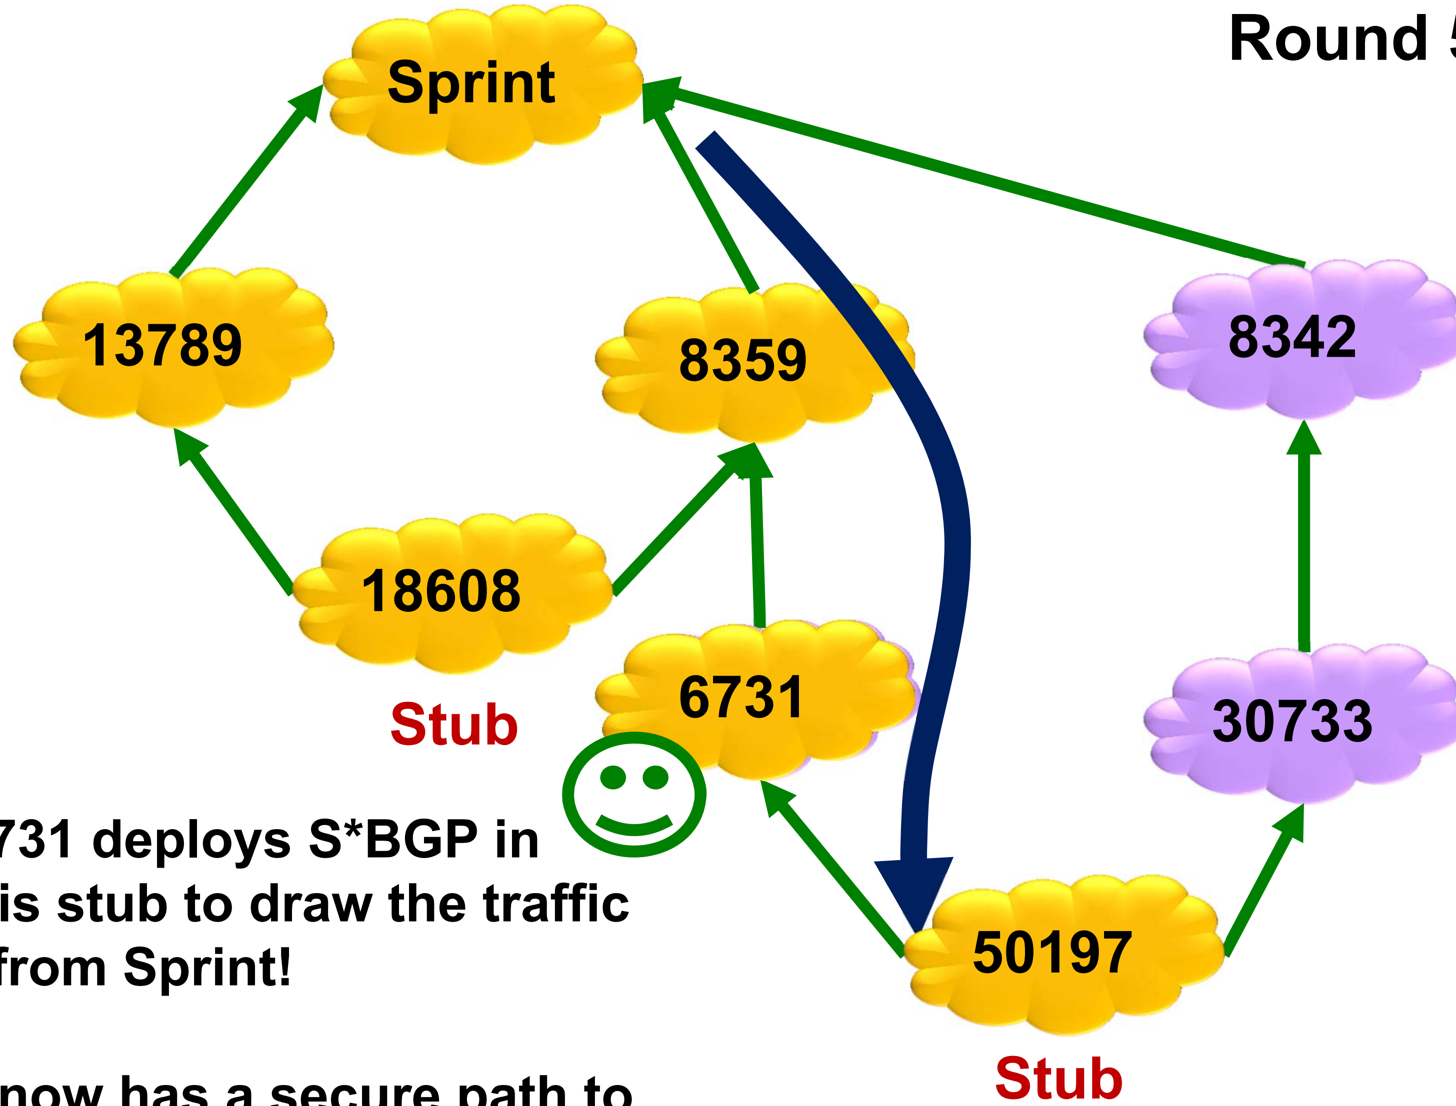

Now AS 6731 deploys S*BGP in himself and his stub to draw the traffic from Sprint!

Notice Sprint now has a secure path to stub 50197 because AS 8359 deployed in the previous round.

Simulation: Market pressure drives deployment (3)

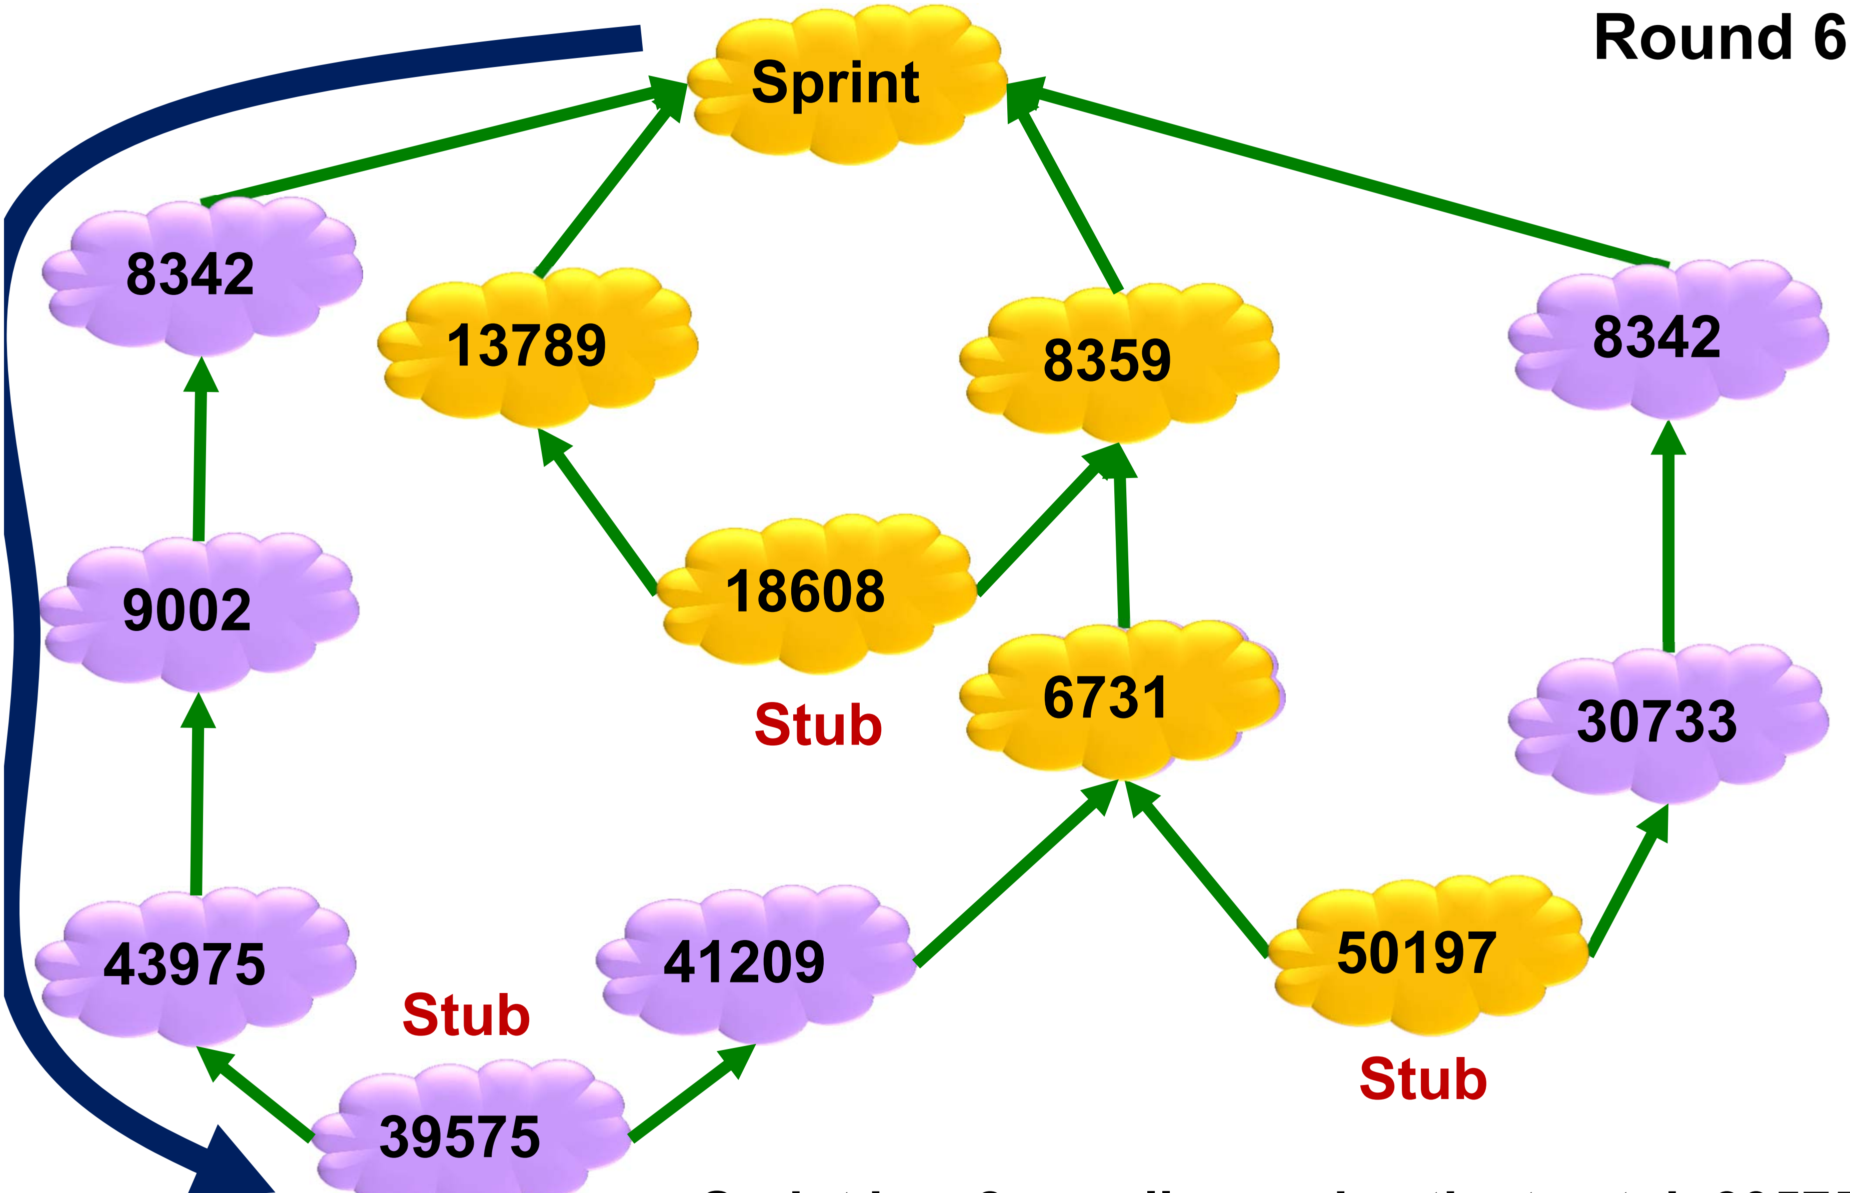

Sprint has 2 equally good paths to stub 39575 `

Simulation: Market pressure drives deployment (3)

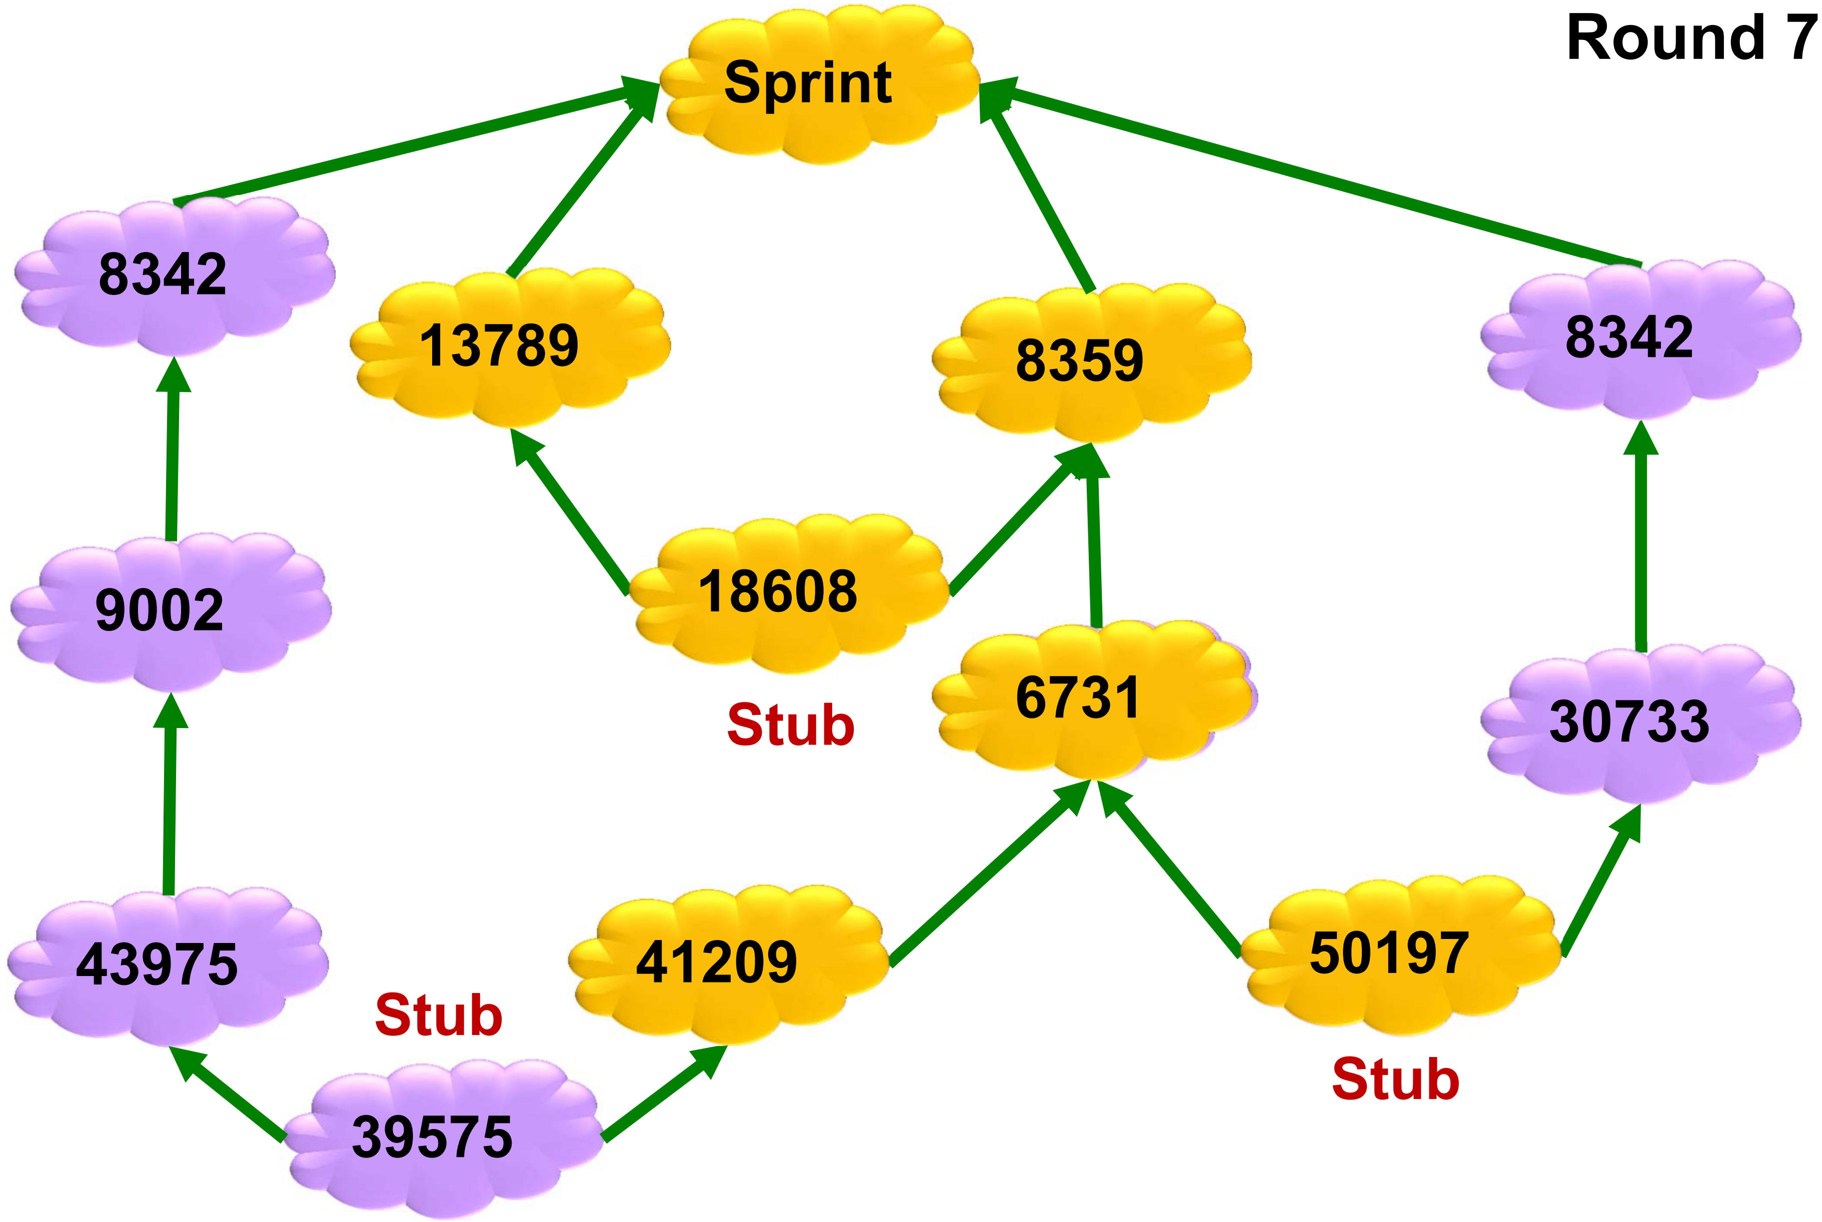

Simulation: Market pressure drives deployment (3)

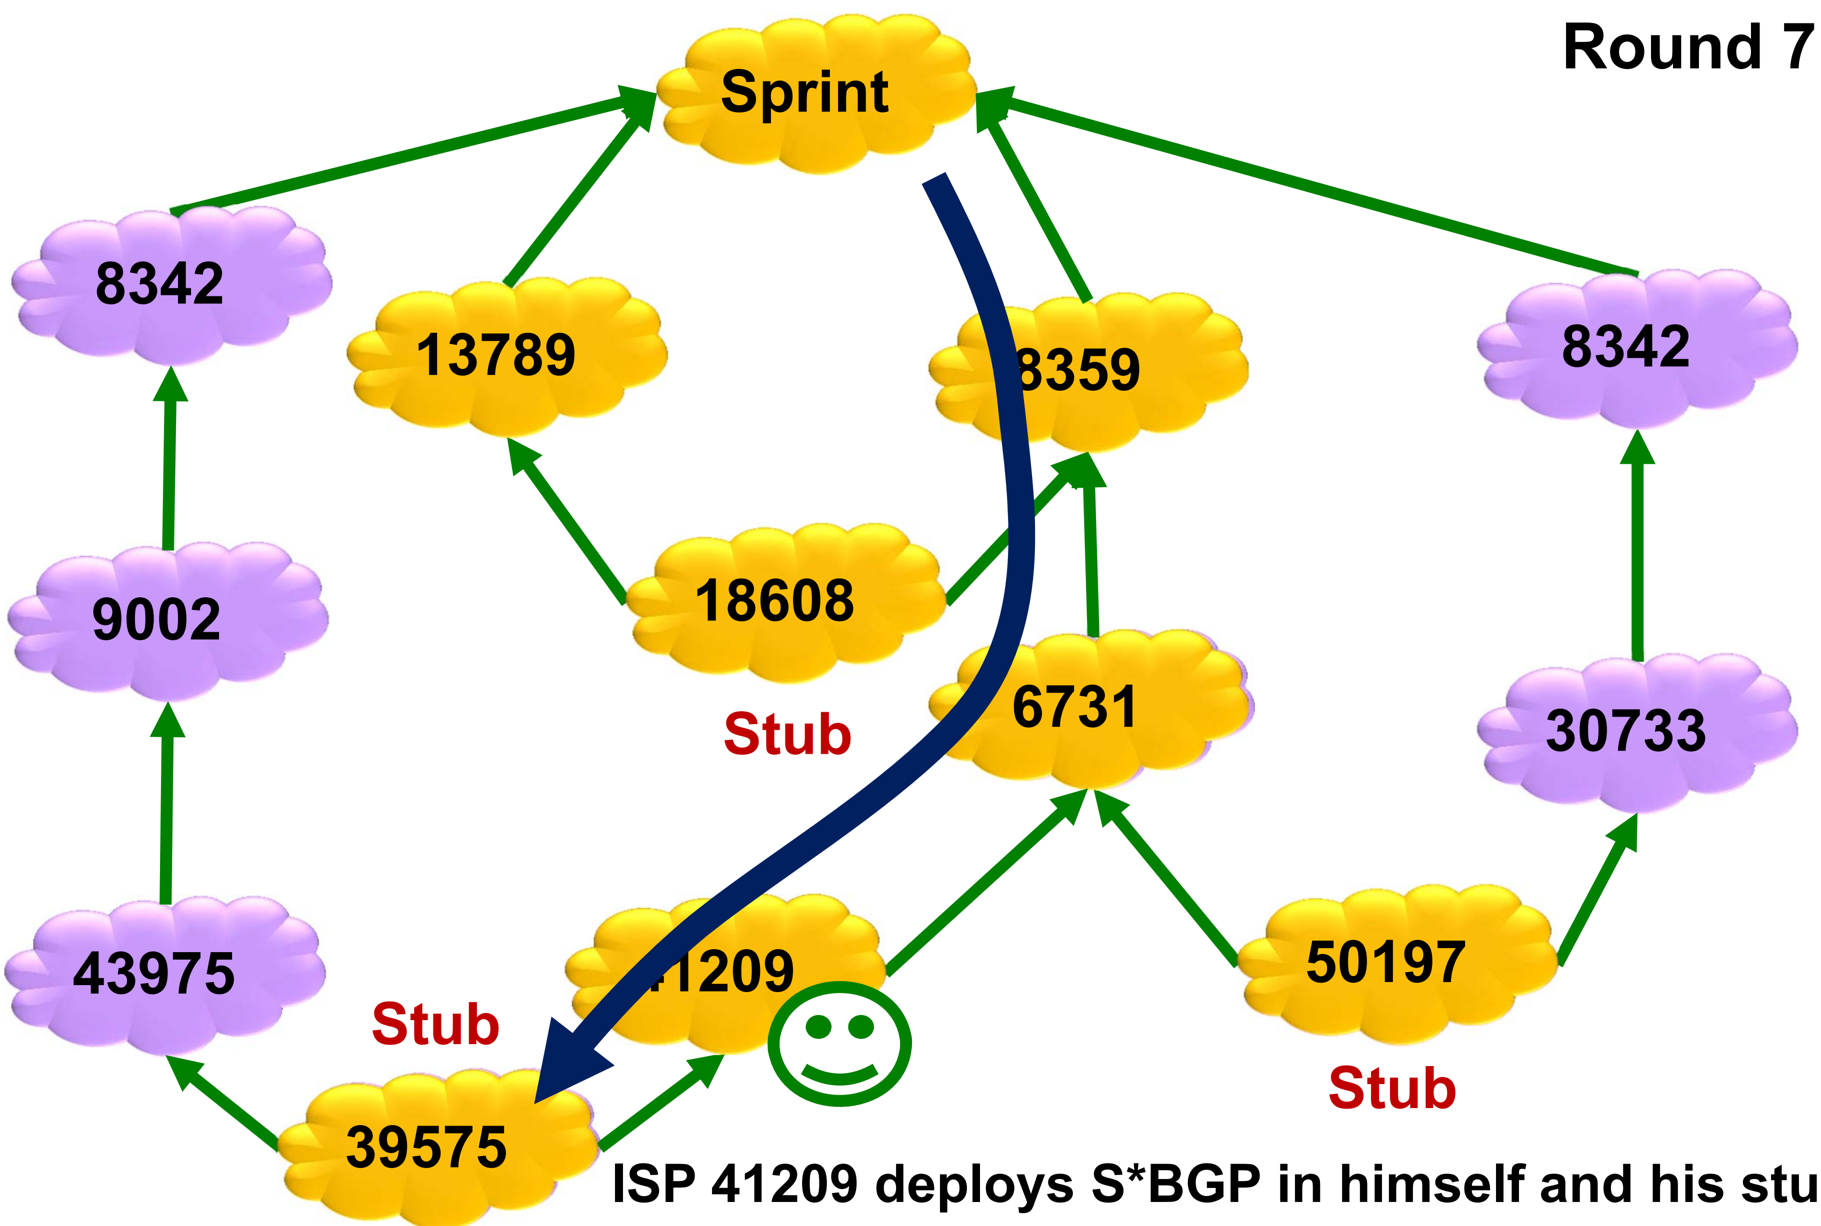

ISP 41209 deploys S*BGP in himself and his stub to get Sprint's traffic.

Changes in traffic volume during deployment (1)

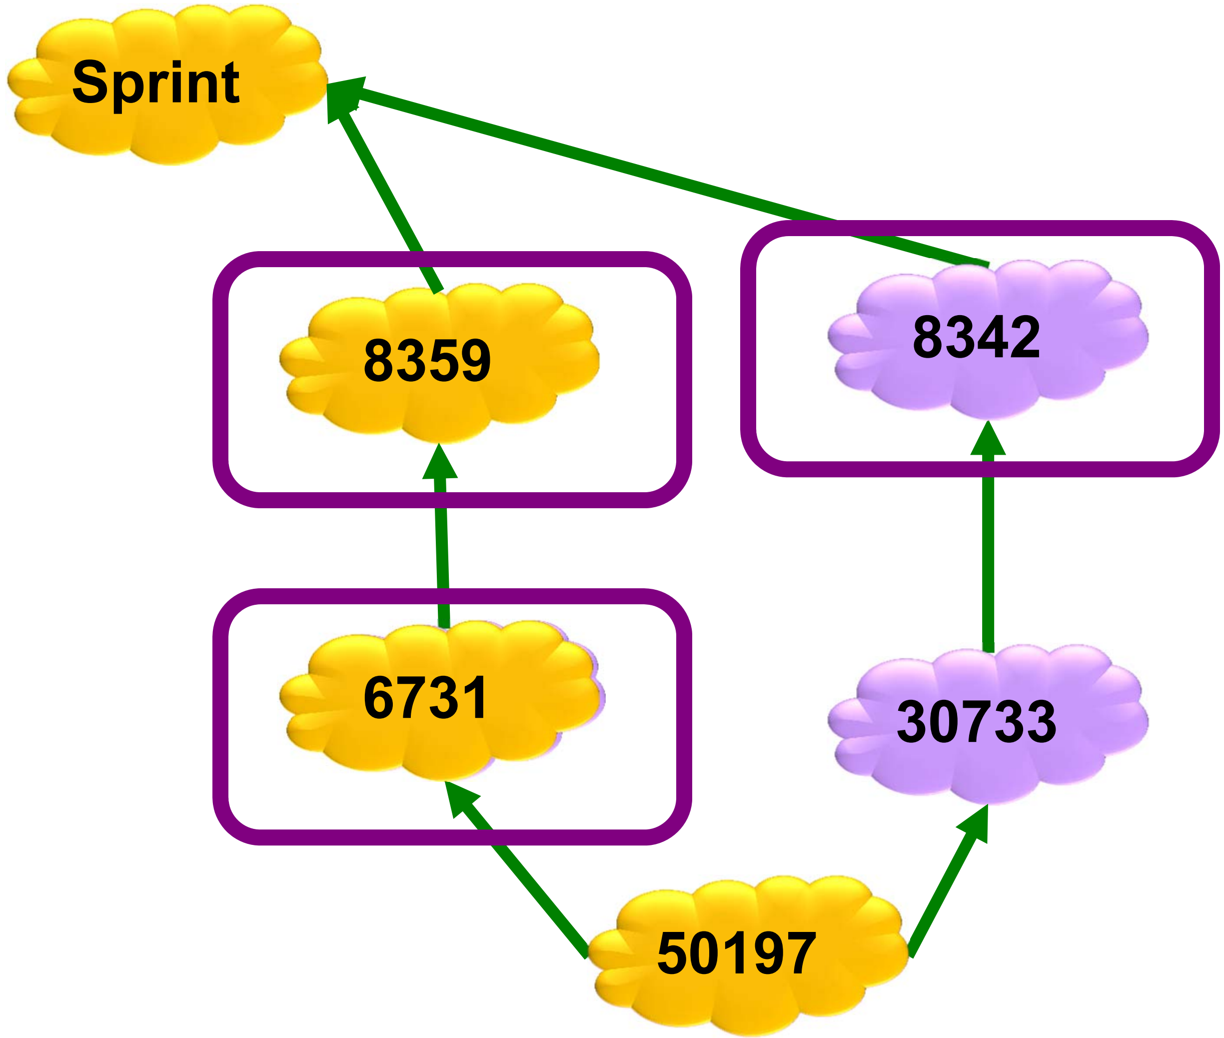

Let's zoom in on the traffic volume at each of these three ISPs…

Changes in traffic volume during deployment (2)

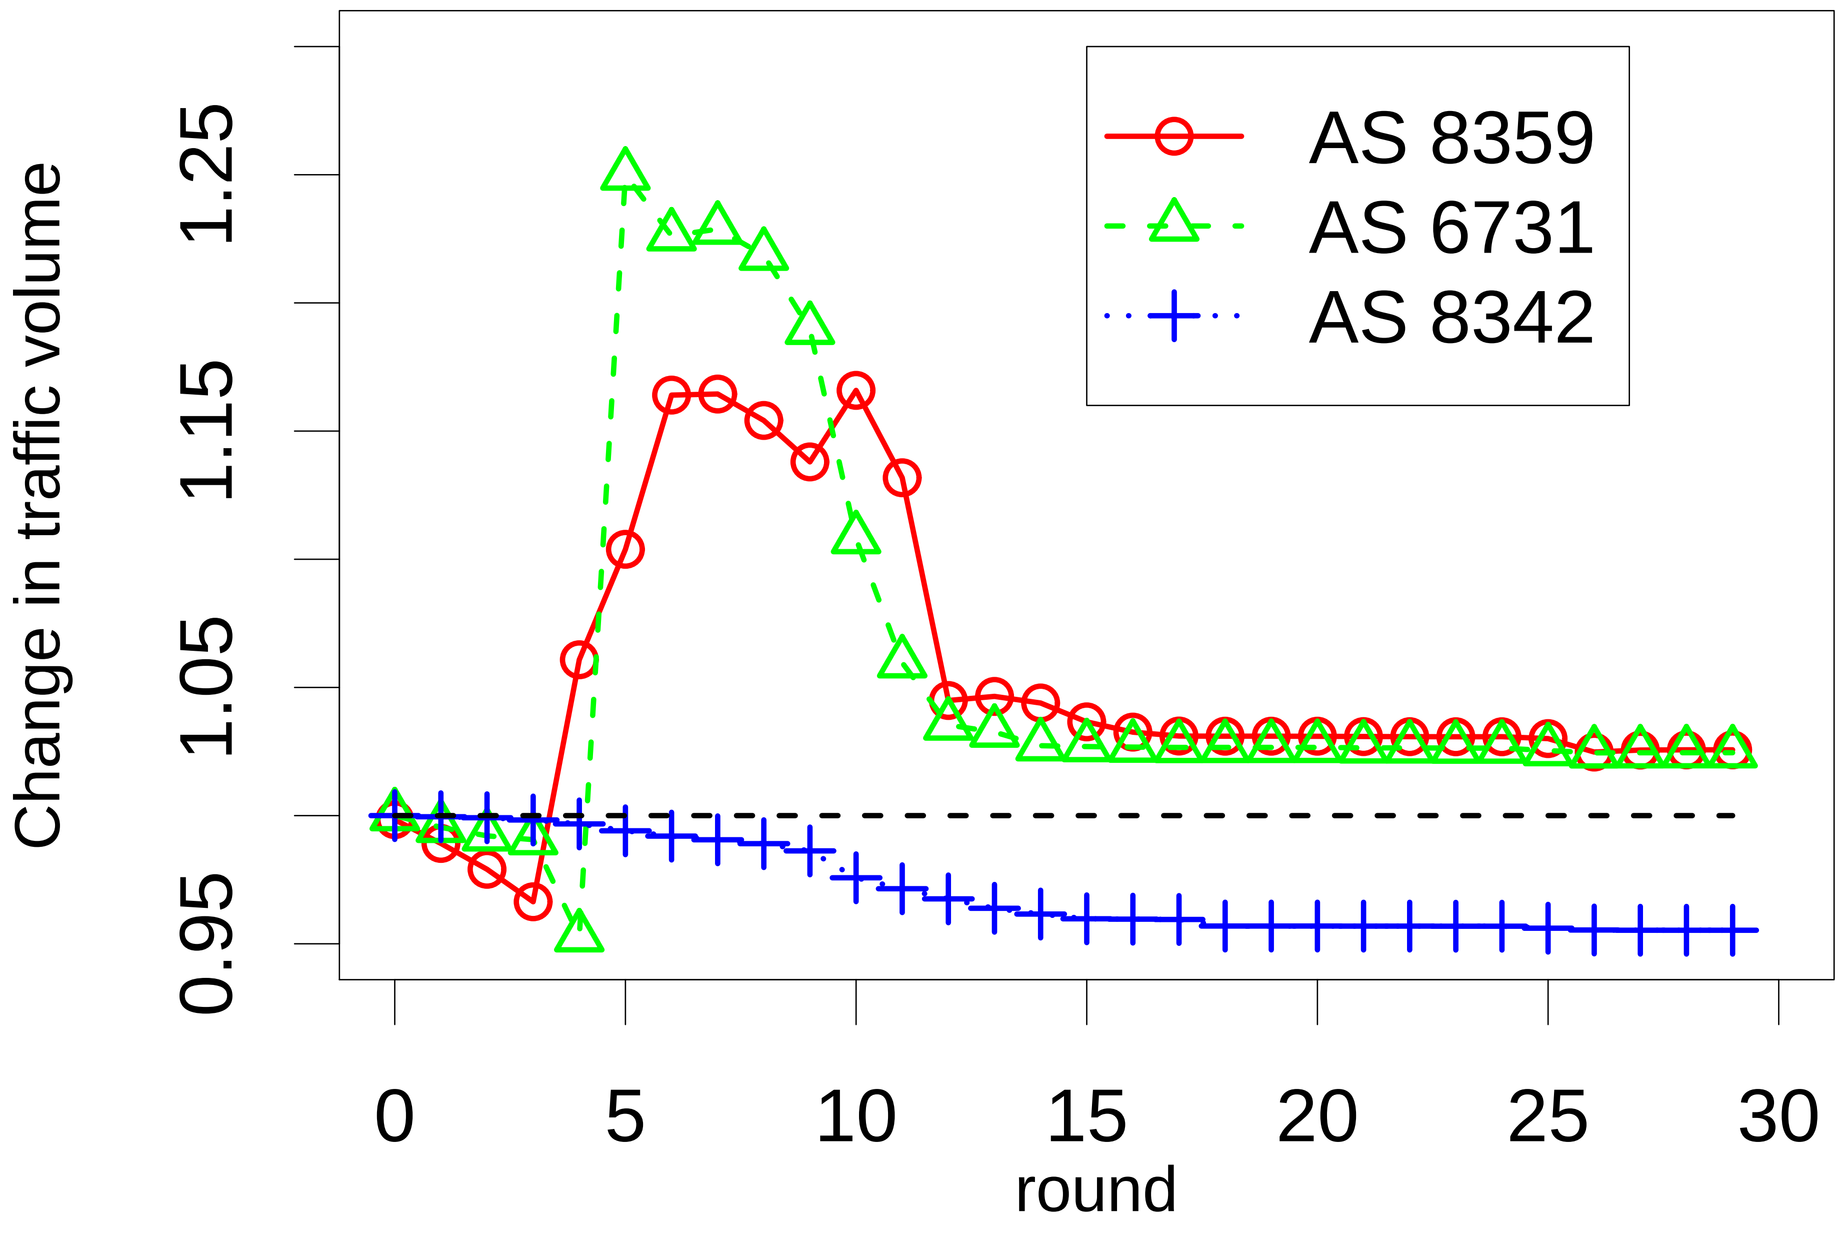

Changes in traffic volume during deployment (3)

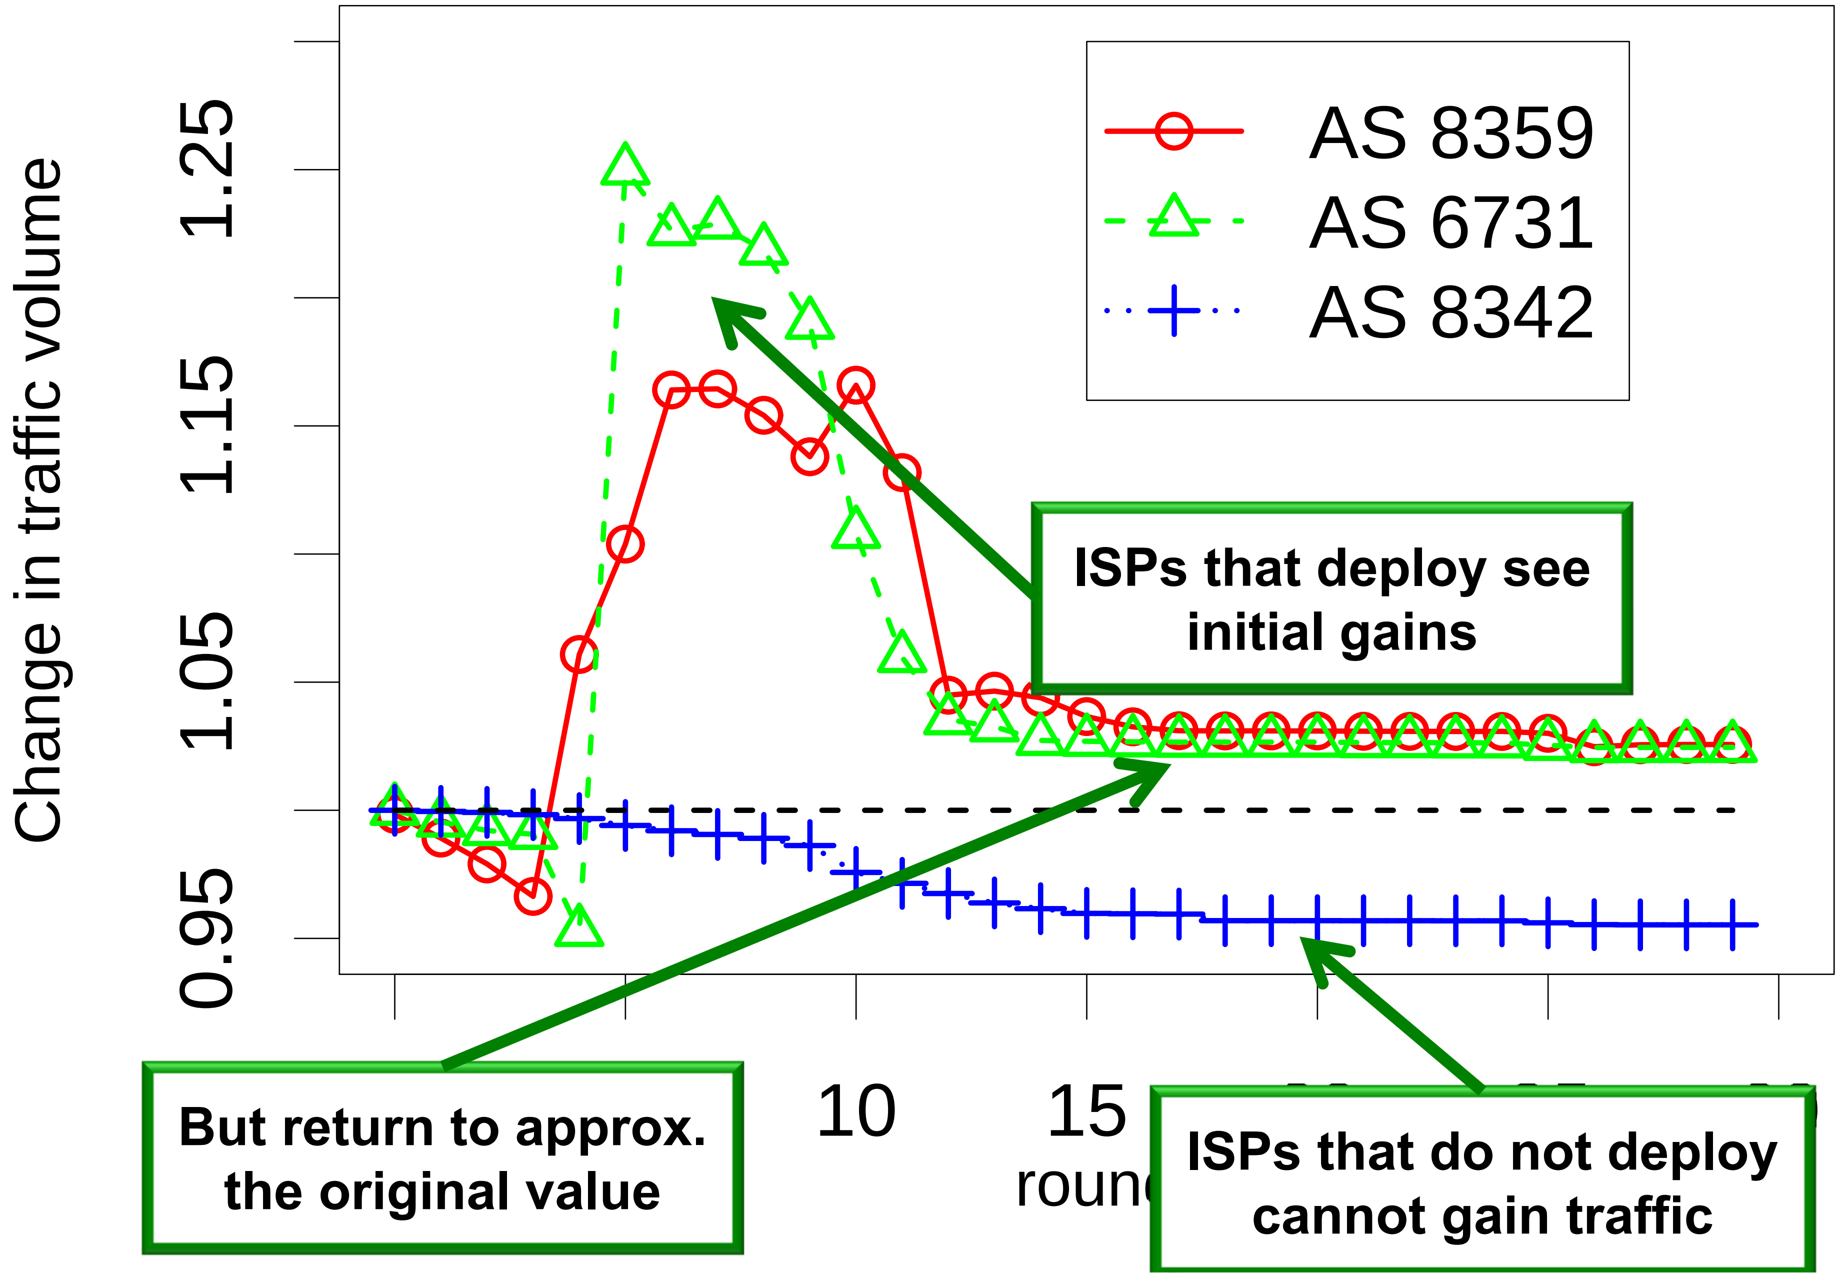

Who should the early adopters be?

At minimum, we need the 5 biggest Tier 1s.

Content providers help, but not as much,
(since they don't have many stub customers)

We ran lots of simulations to figure this out.
See our tech report.

So who should the early adopters be?

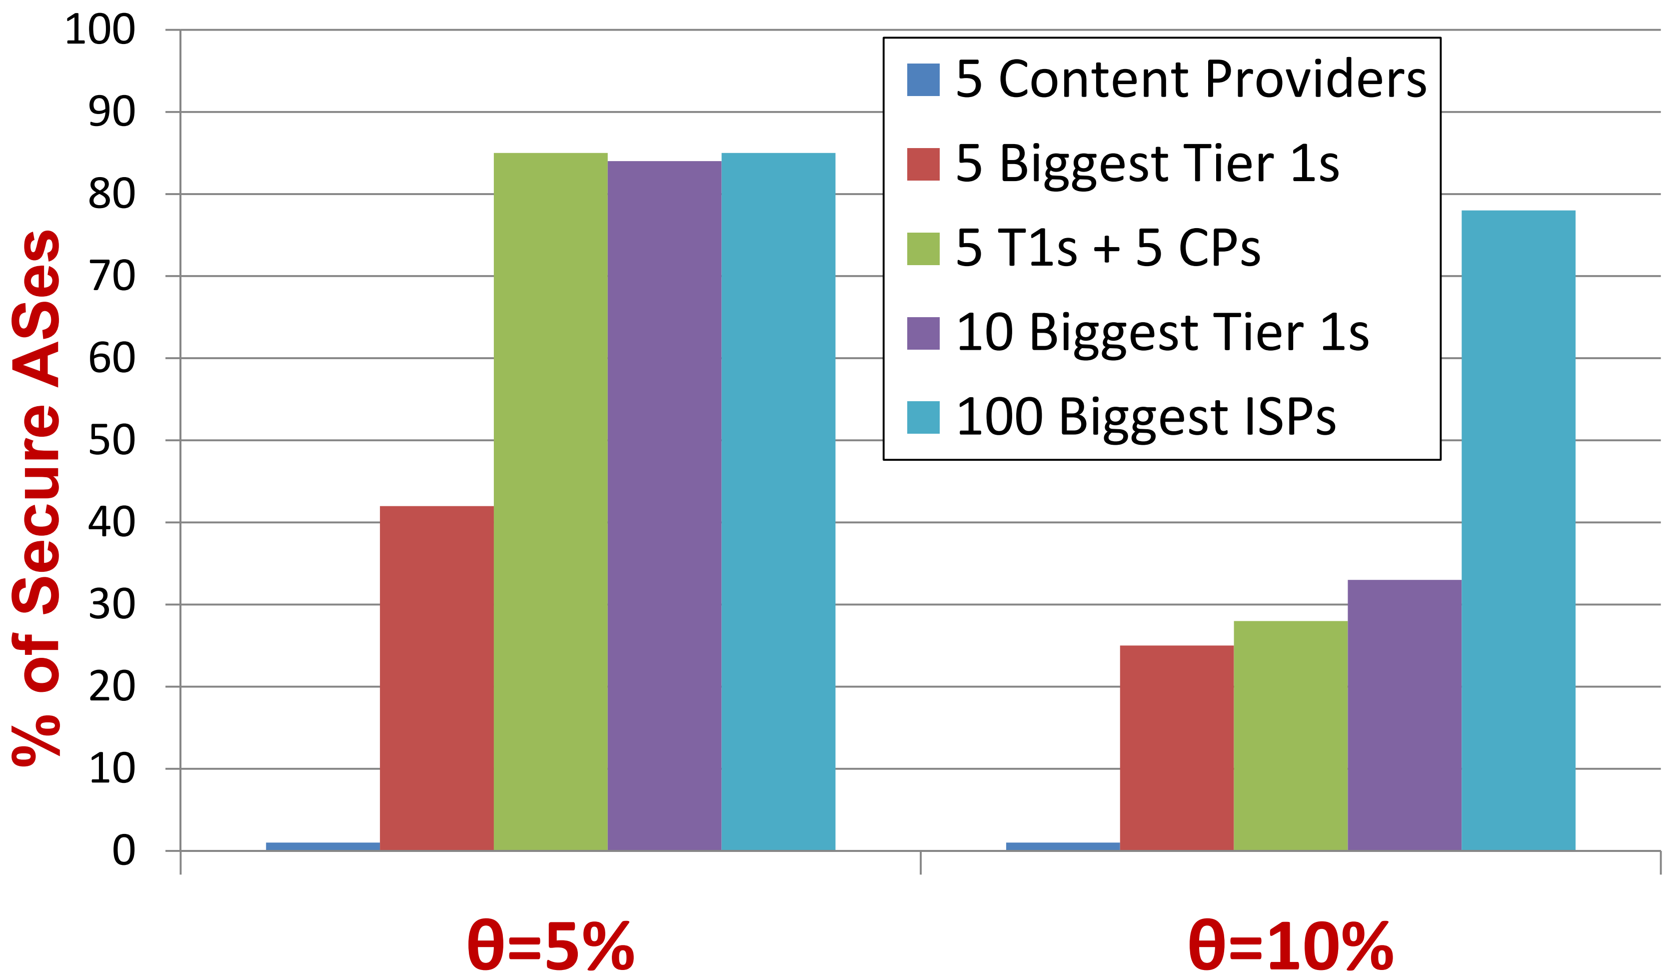

In ISPs are willing to re-invest θ% of new revenue from increases in attracted traffic in S*BGP, then only a few early adopters are enough to drive (almost) global deployment.

Observations and Recommendations

To improve security, S*BGP should impact route selection

1. Thus it has an impact on **traffic engineering**.

2. But it's also an opportunity to offer **differentiated services**

   … and attract customers away from your competitors

   … so that deployment at your ISP "pays for itself".

Where should gov't funding and regulation go?

1. **Subsidize early adopters**: Tier 1s / content providers
2. **Subsidize ISPs** that upgrade stubs to **simplex S*BGP**
   - Crucial since **85%** of ASes are stubs
   - ISPs, it's really important you involve your customers.

This work is not predictive!

Instead, our goal was to capture key issues affecting deployment.

Thanks!

This work will also appear at SIGCOMM'11

Detailed results are in our tech report:
http://www.cs.bu.edu/~goldbe/papers/sbgpTrans.html

Also, download our interactive results browser console app at the above url & browse our full simulation results.

BOSTON
UNIVERSITY

http://www.cs.bu.edu/~goldbe
goldbe@cs.bu.edu

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**

85% of the Internet's ASes are stubs.

A stub never transits traffic!

- Thus, it only sends BGP messages
  … for **its own prefixes,** and for
  … paths that are **exactly 0 hops long.**

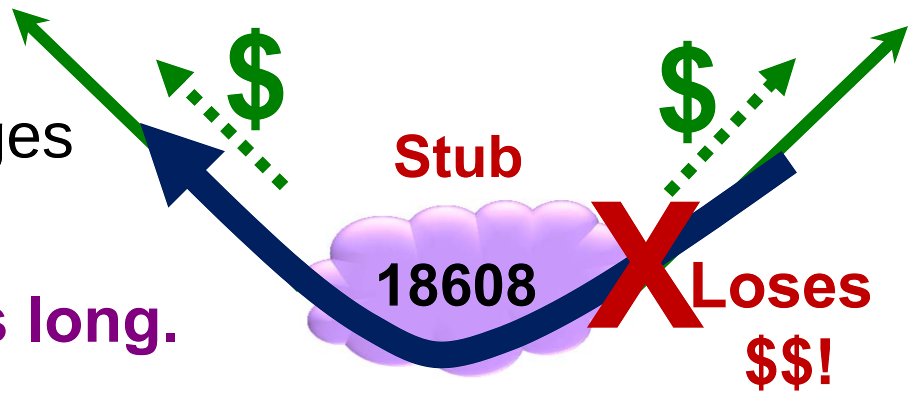

2 options for deploying S*BGP in stubs:

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.

Tiebreak Sets: The Source of Competition (1)

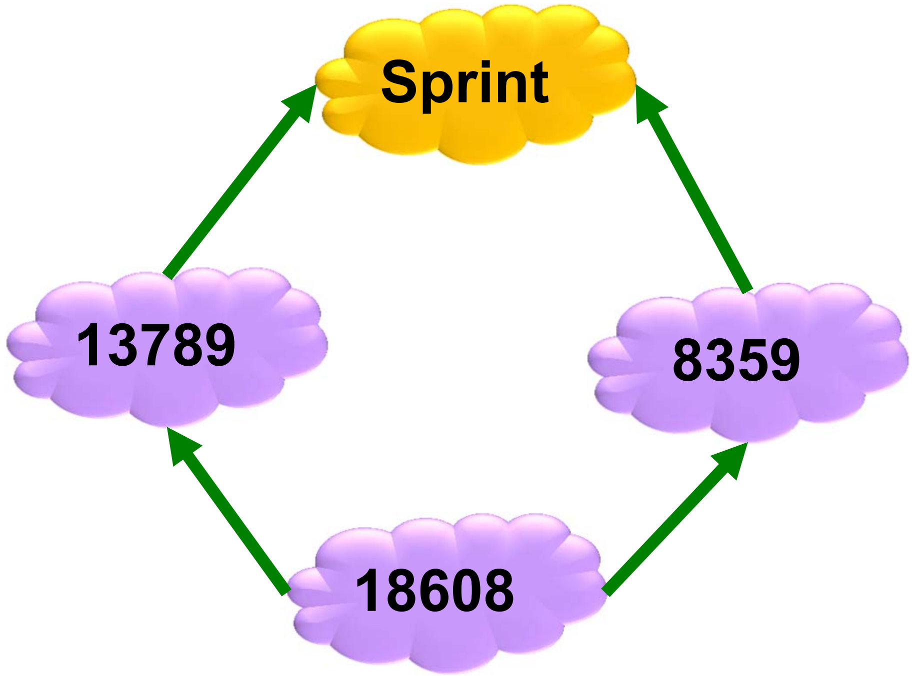

Sprint's tiebreak set to destination **AS18608** is

{AS 13789, AS 8357}

Thus, these two ISPs compete for traffic!

Tiebreak Sets: The Source of Competition (2)

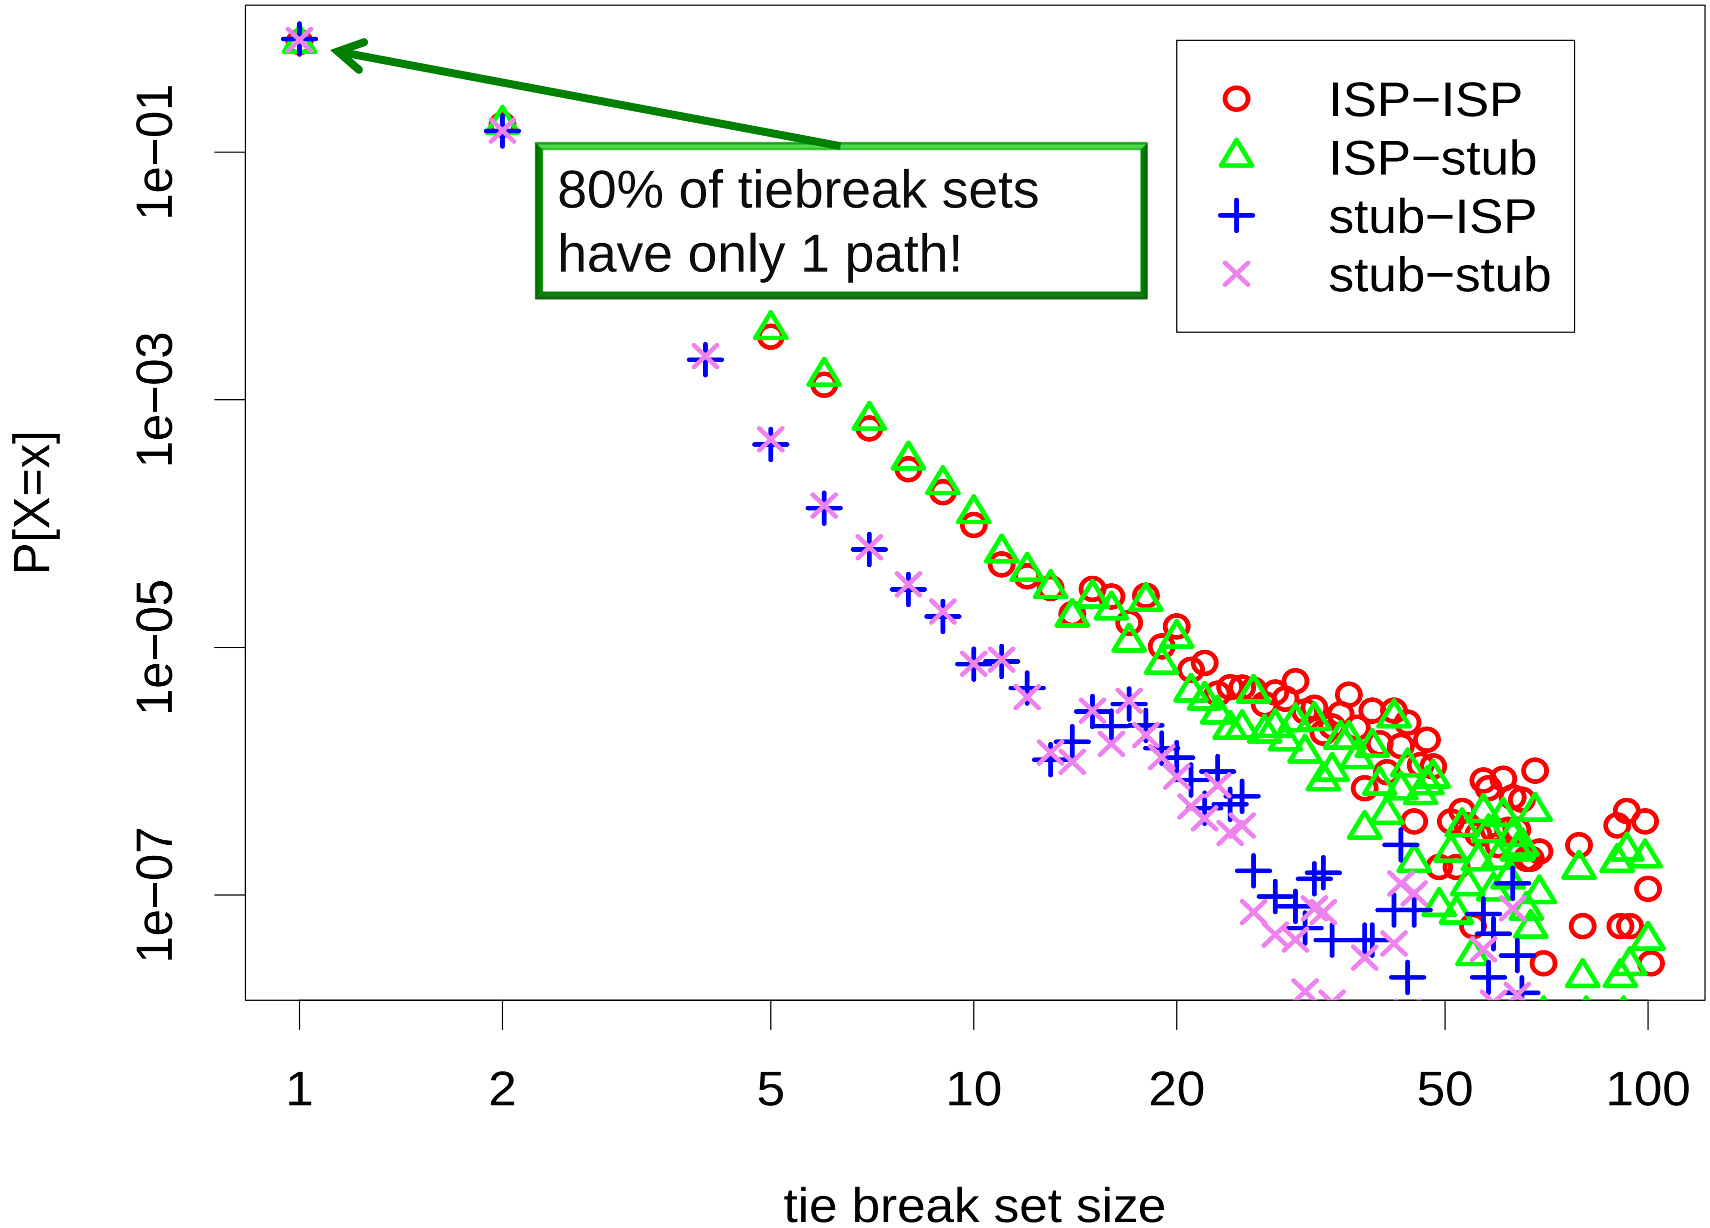

Tiebreak Sets: The Source of Competition (3)

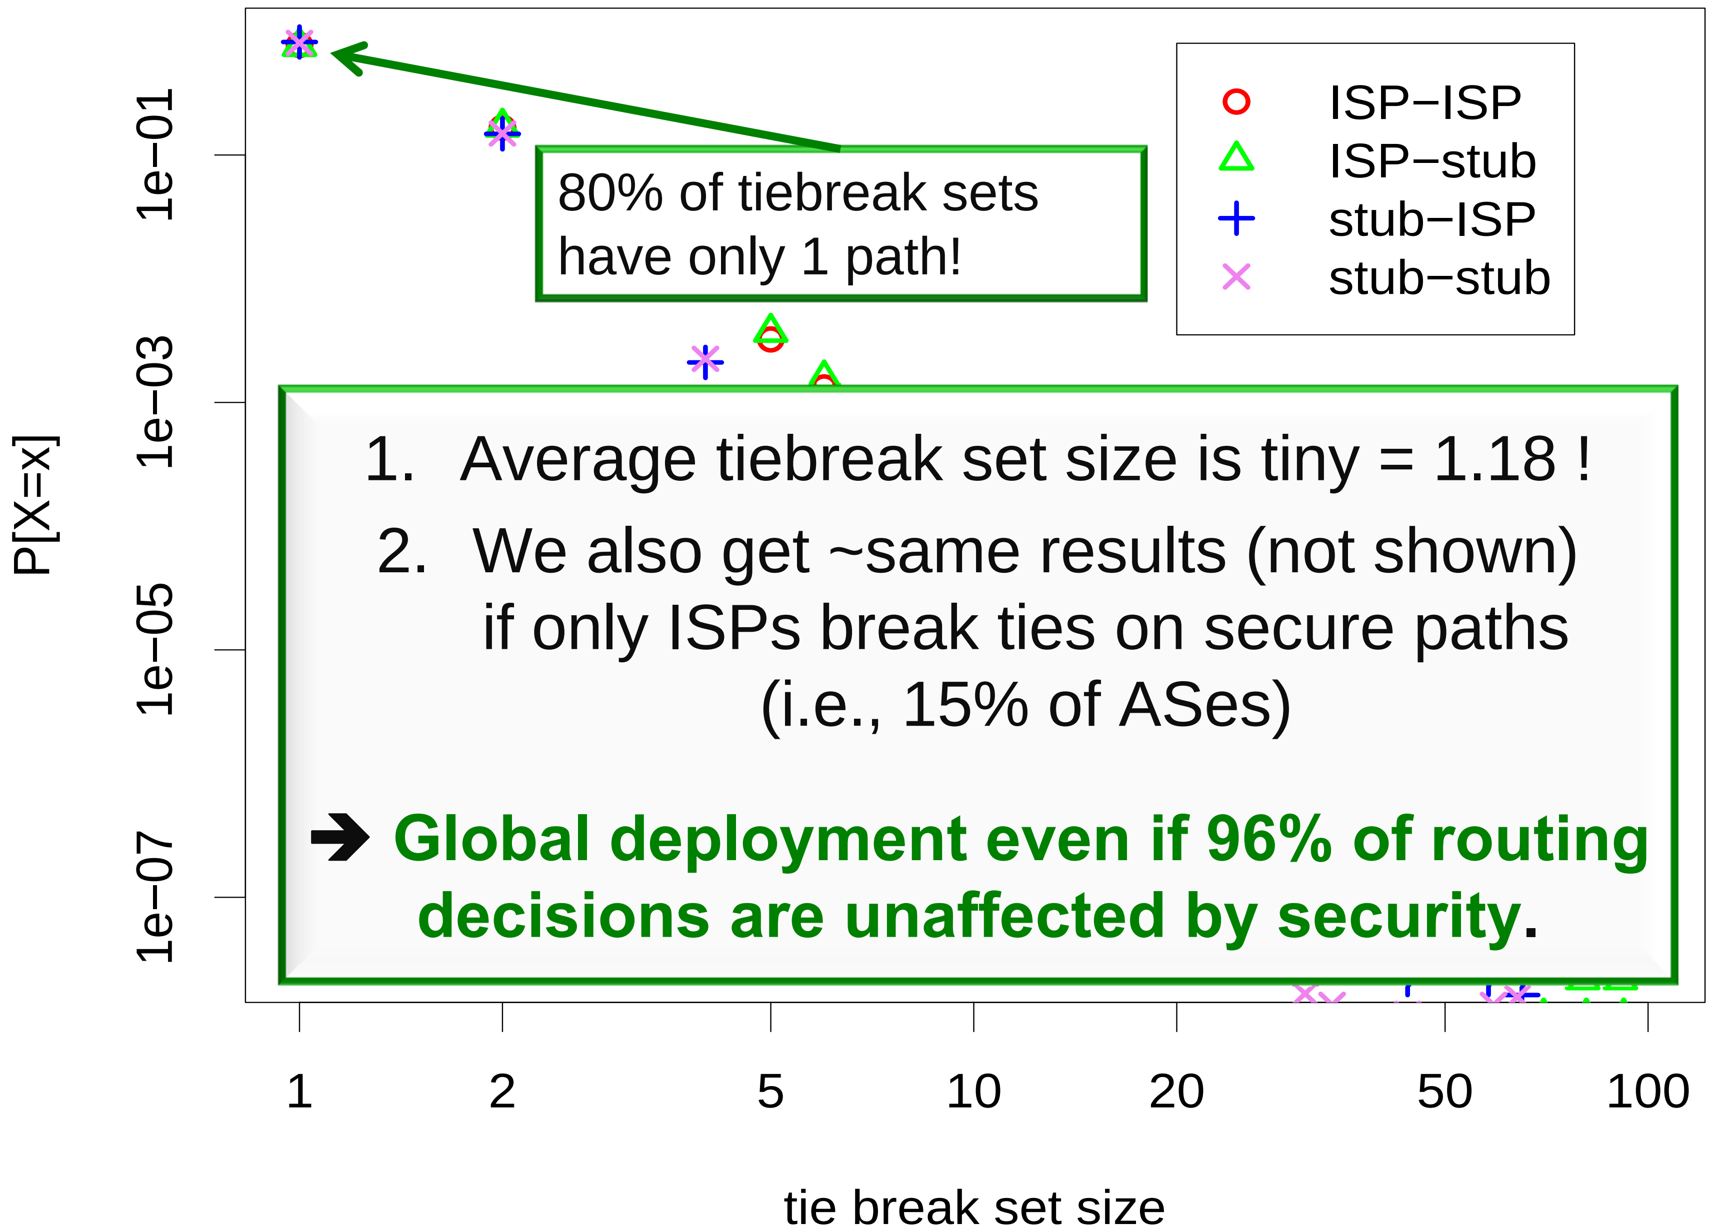

1. Average tiebreak set size is tiny = 1.18 !
2. We also get ~same results (not shown) if only ISPs break ties on secure paths (i.e., 15% of ASes)

➔ **Global deployment even if 96% of routing decisions are unaffected by security.**

So who should be the early adopters?

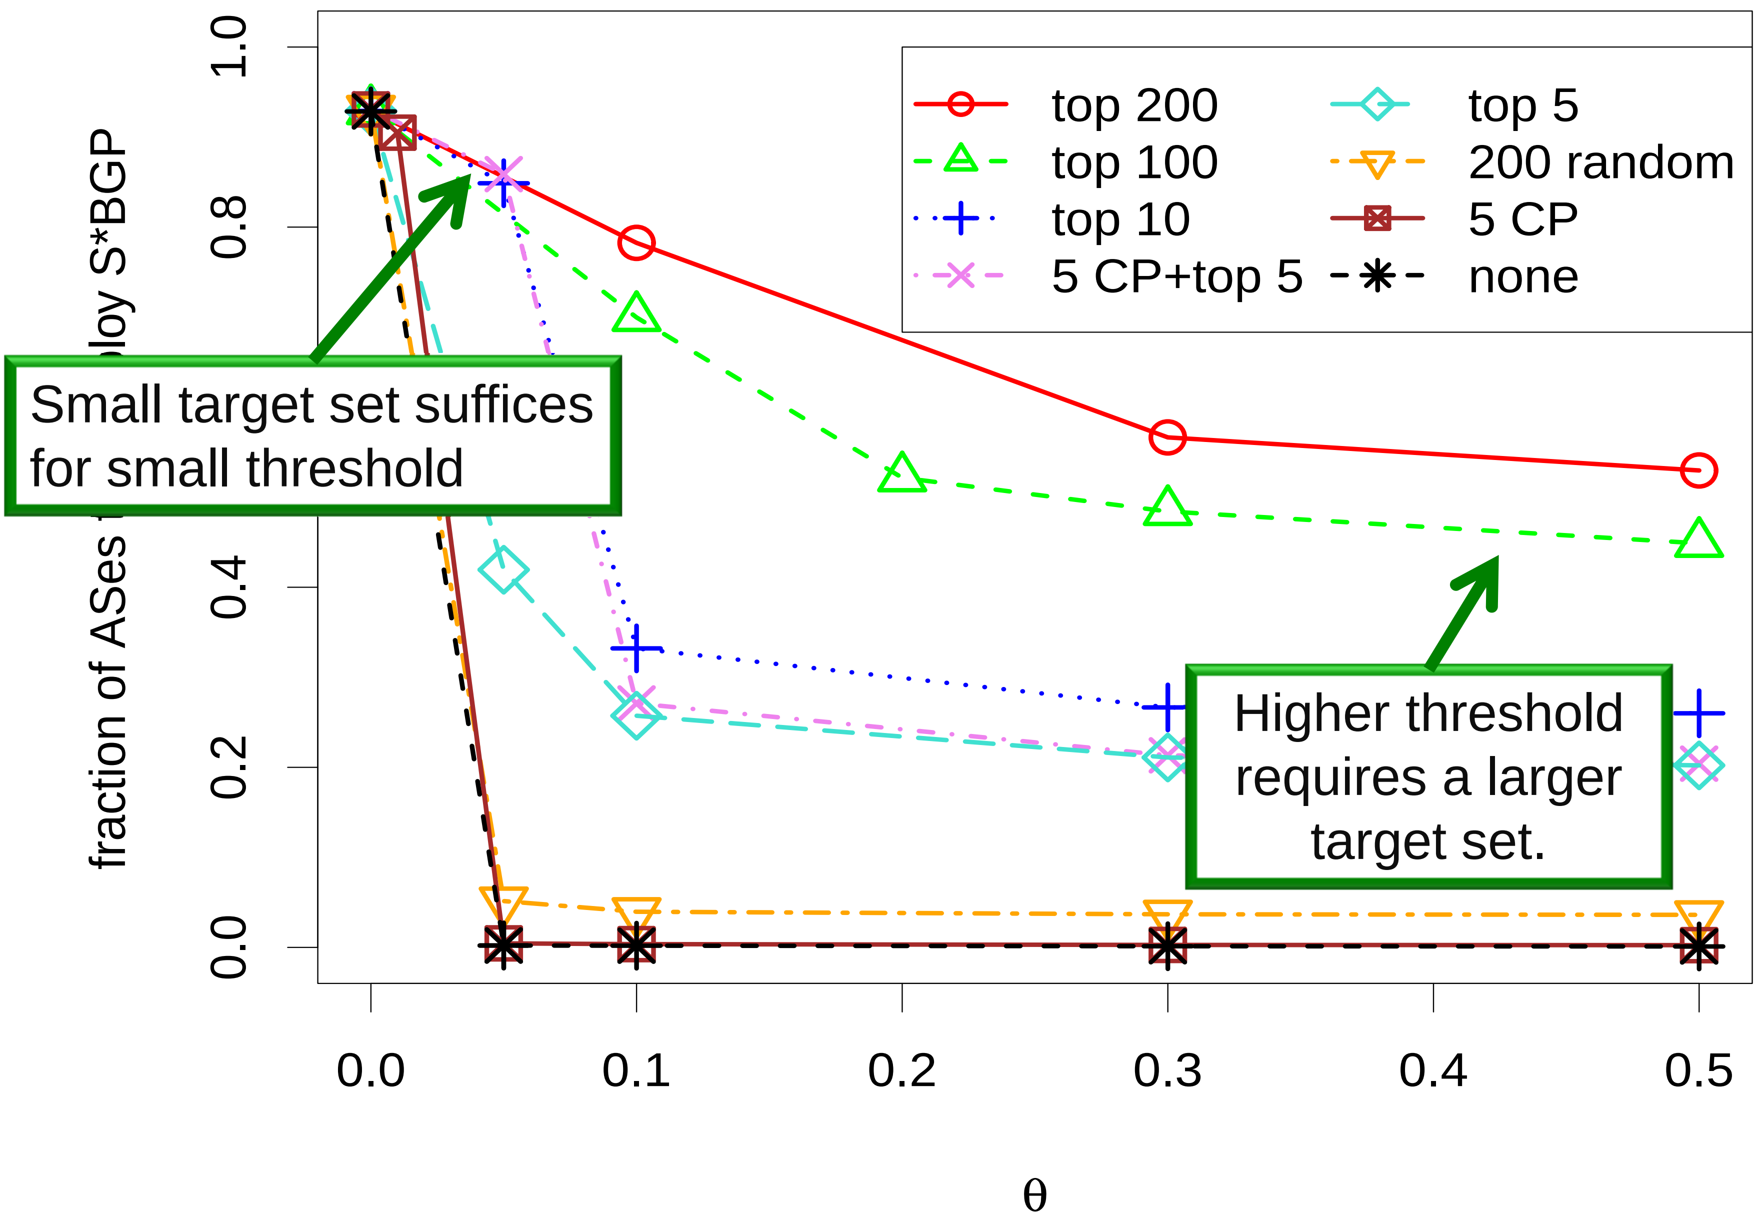

Simplex S*BGP vs. Market-pressure