Improving Routing Security for Global Research and Education

Matthew Luckie
CAIDA

Steven Wallace
Internet2

Karl Newell
Internet2

Jeff Bartig
Internet2

James Deaton
Internet2

k claffy
CAIDA

Part of ROOTBEER Project.
Routing Operations Observational Technology:
Building to Enable Education and Research

NSF CICI OAC-2530871

https://www.caida.org/funding/cici-rootbeer/

Research and Education (R&E) Networking: Primer

- **Specialized R&E networking infrastructure bridges a gap:**
    - R&E institutions need to exchange large volumes of scientific data globally
    - Impractical using commercial CDNs
- Primary focus of R&E networking is on high-speed low-latency connectivity between members
- **Members may have to arrange their own commodity transit**

Global R&E Networking Infrastructure Interconnects 104 Countries

~2,700 ASs
~12,000 IPv4 Prefixes

R&E Networking: Primer

Which route does AS B choose?

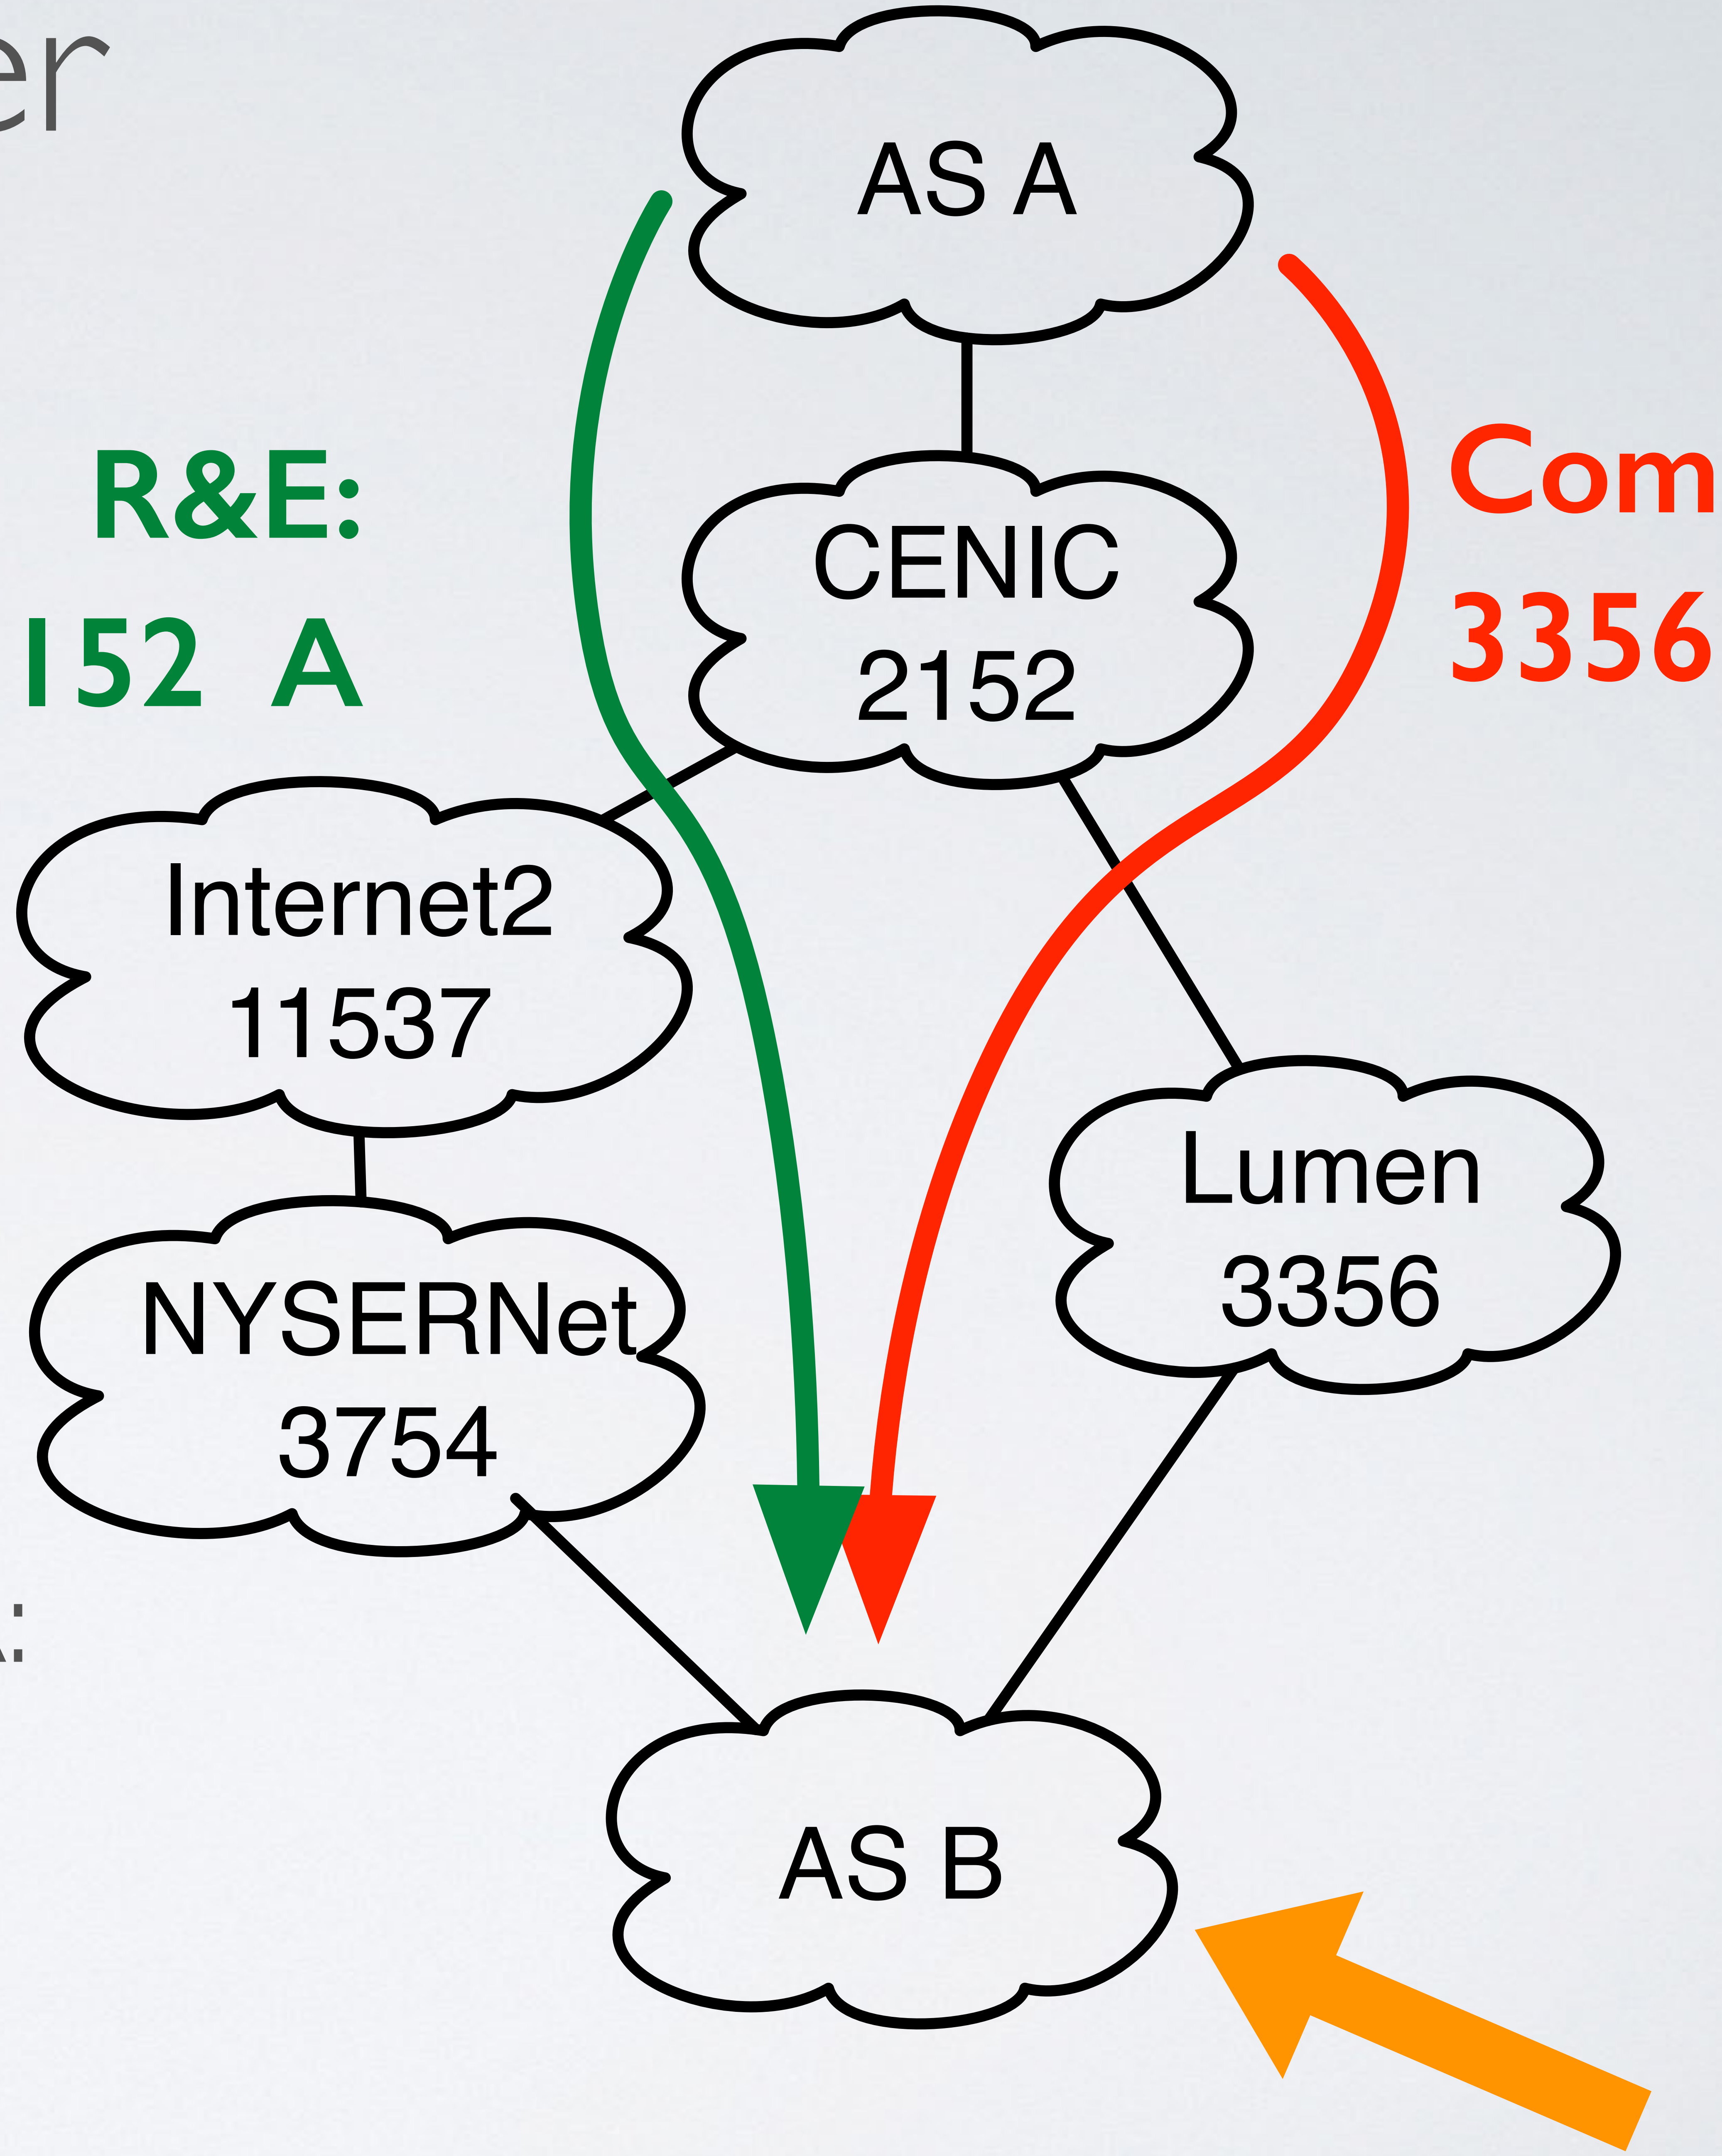

R&E:

3754 11537 2152 A

Commodity:

3356 2152 A

AS B may receive two routes to AS A:

commodity via 3356 (Lumen)

R&E via 3754 (NYSERNet)

R&E Networking: Primer

Which route does AS B choose?

R&E:

3754 11537 2152 A

Commodity:

3356 2152 A

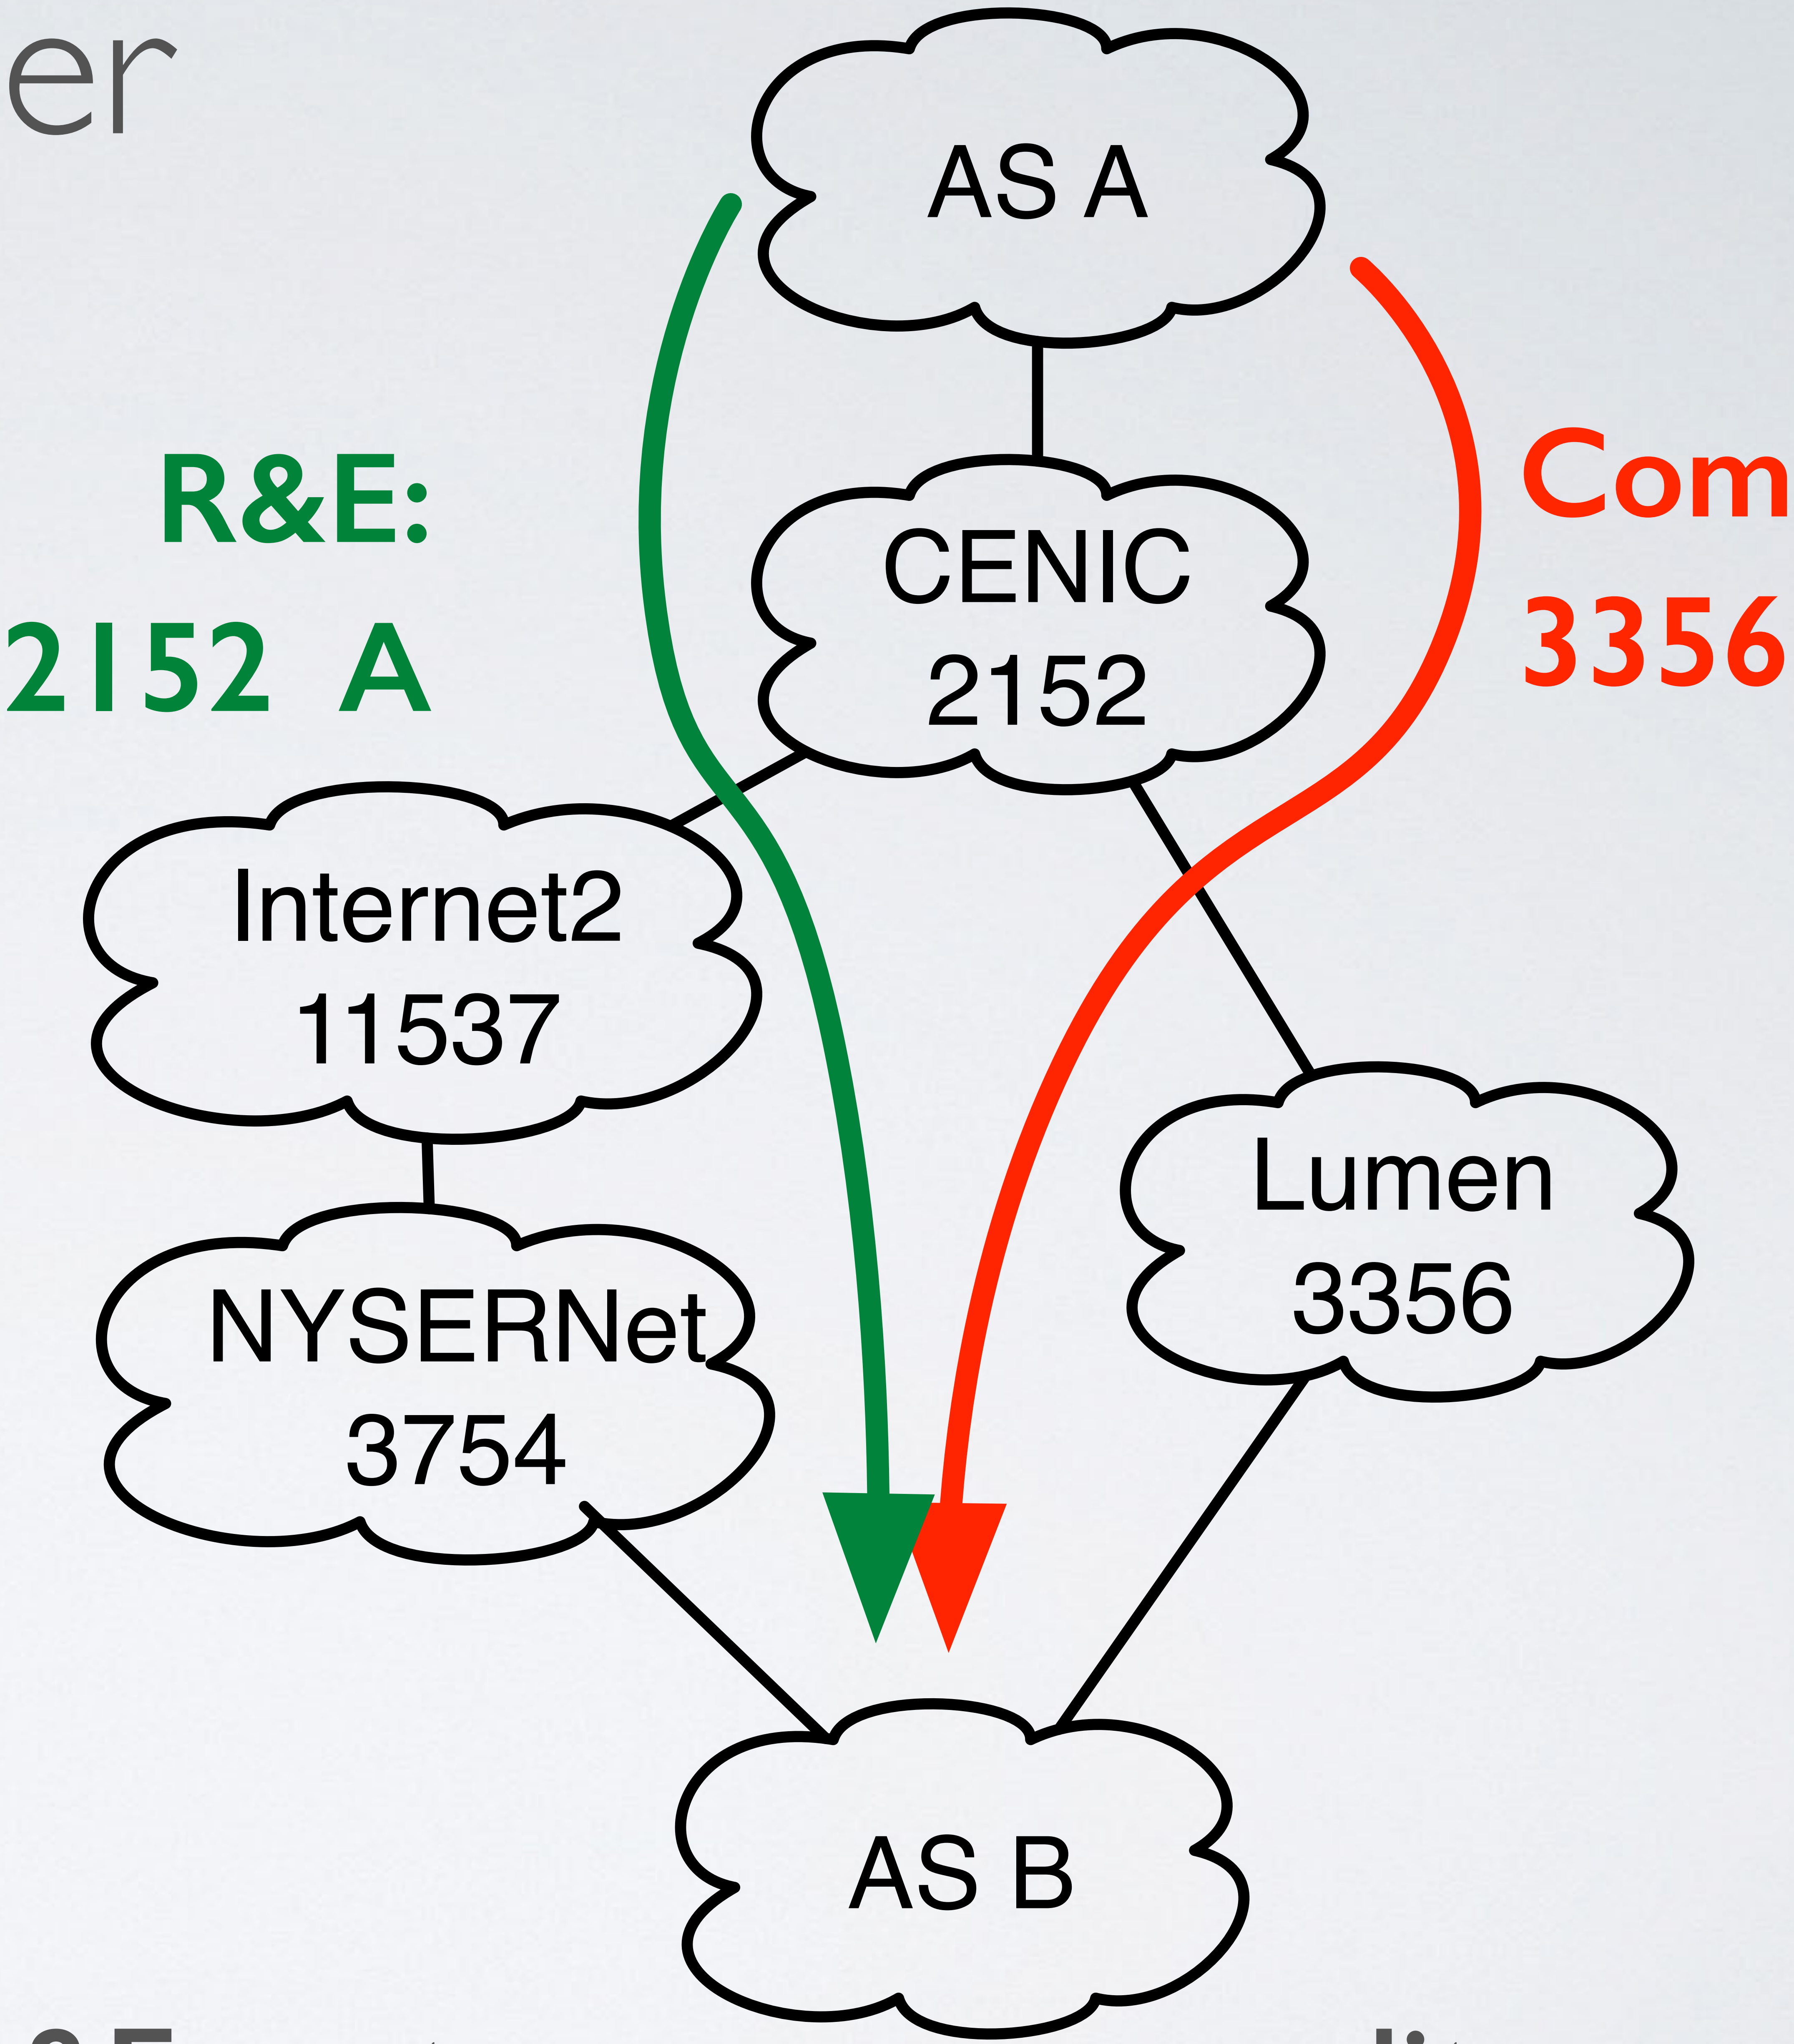

To deterministically select R&E routes, AS B has to set a higher localpref on R&E routes than commodity routes

Do R&E networks prefer R&E routes over commodity routes?

Which route does AS B choose?

- BGP is an information hiding protocol
- BGP does not provide visibility into decisions downstream of a view
- AS A can’t tell which route AS B will select

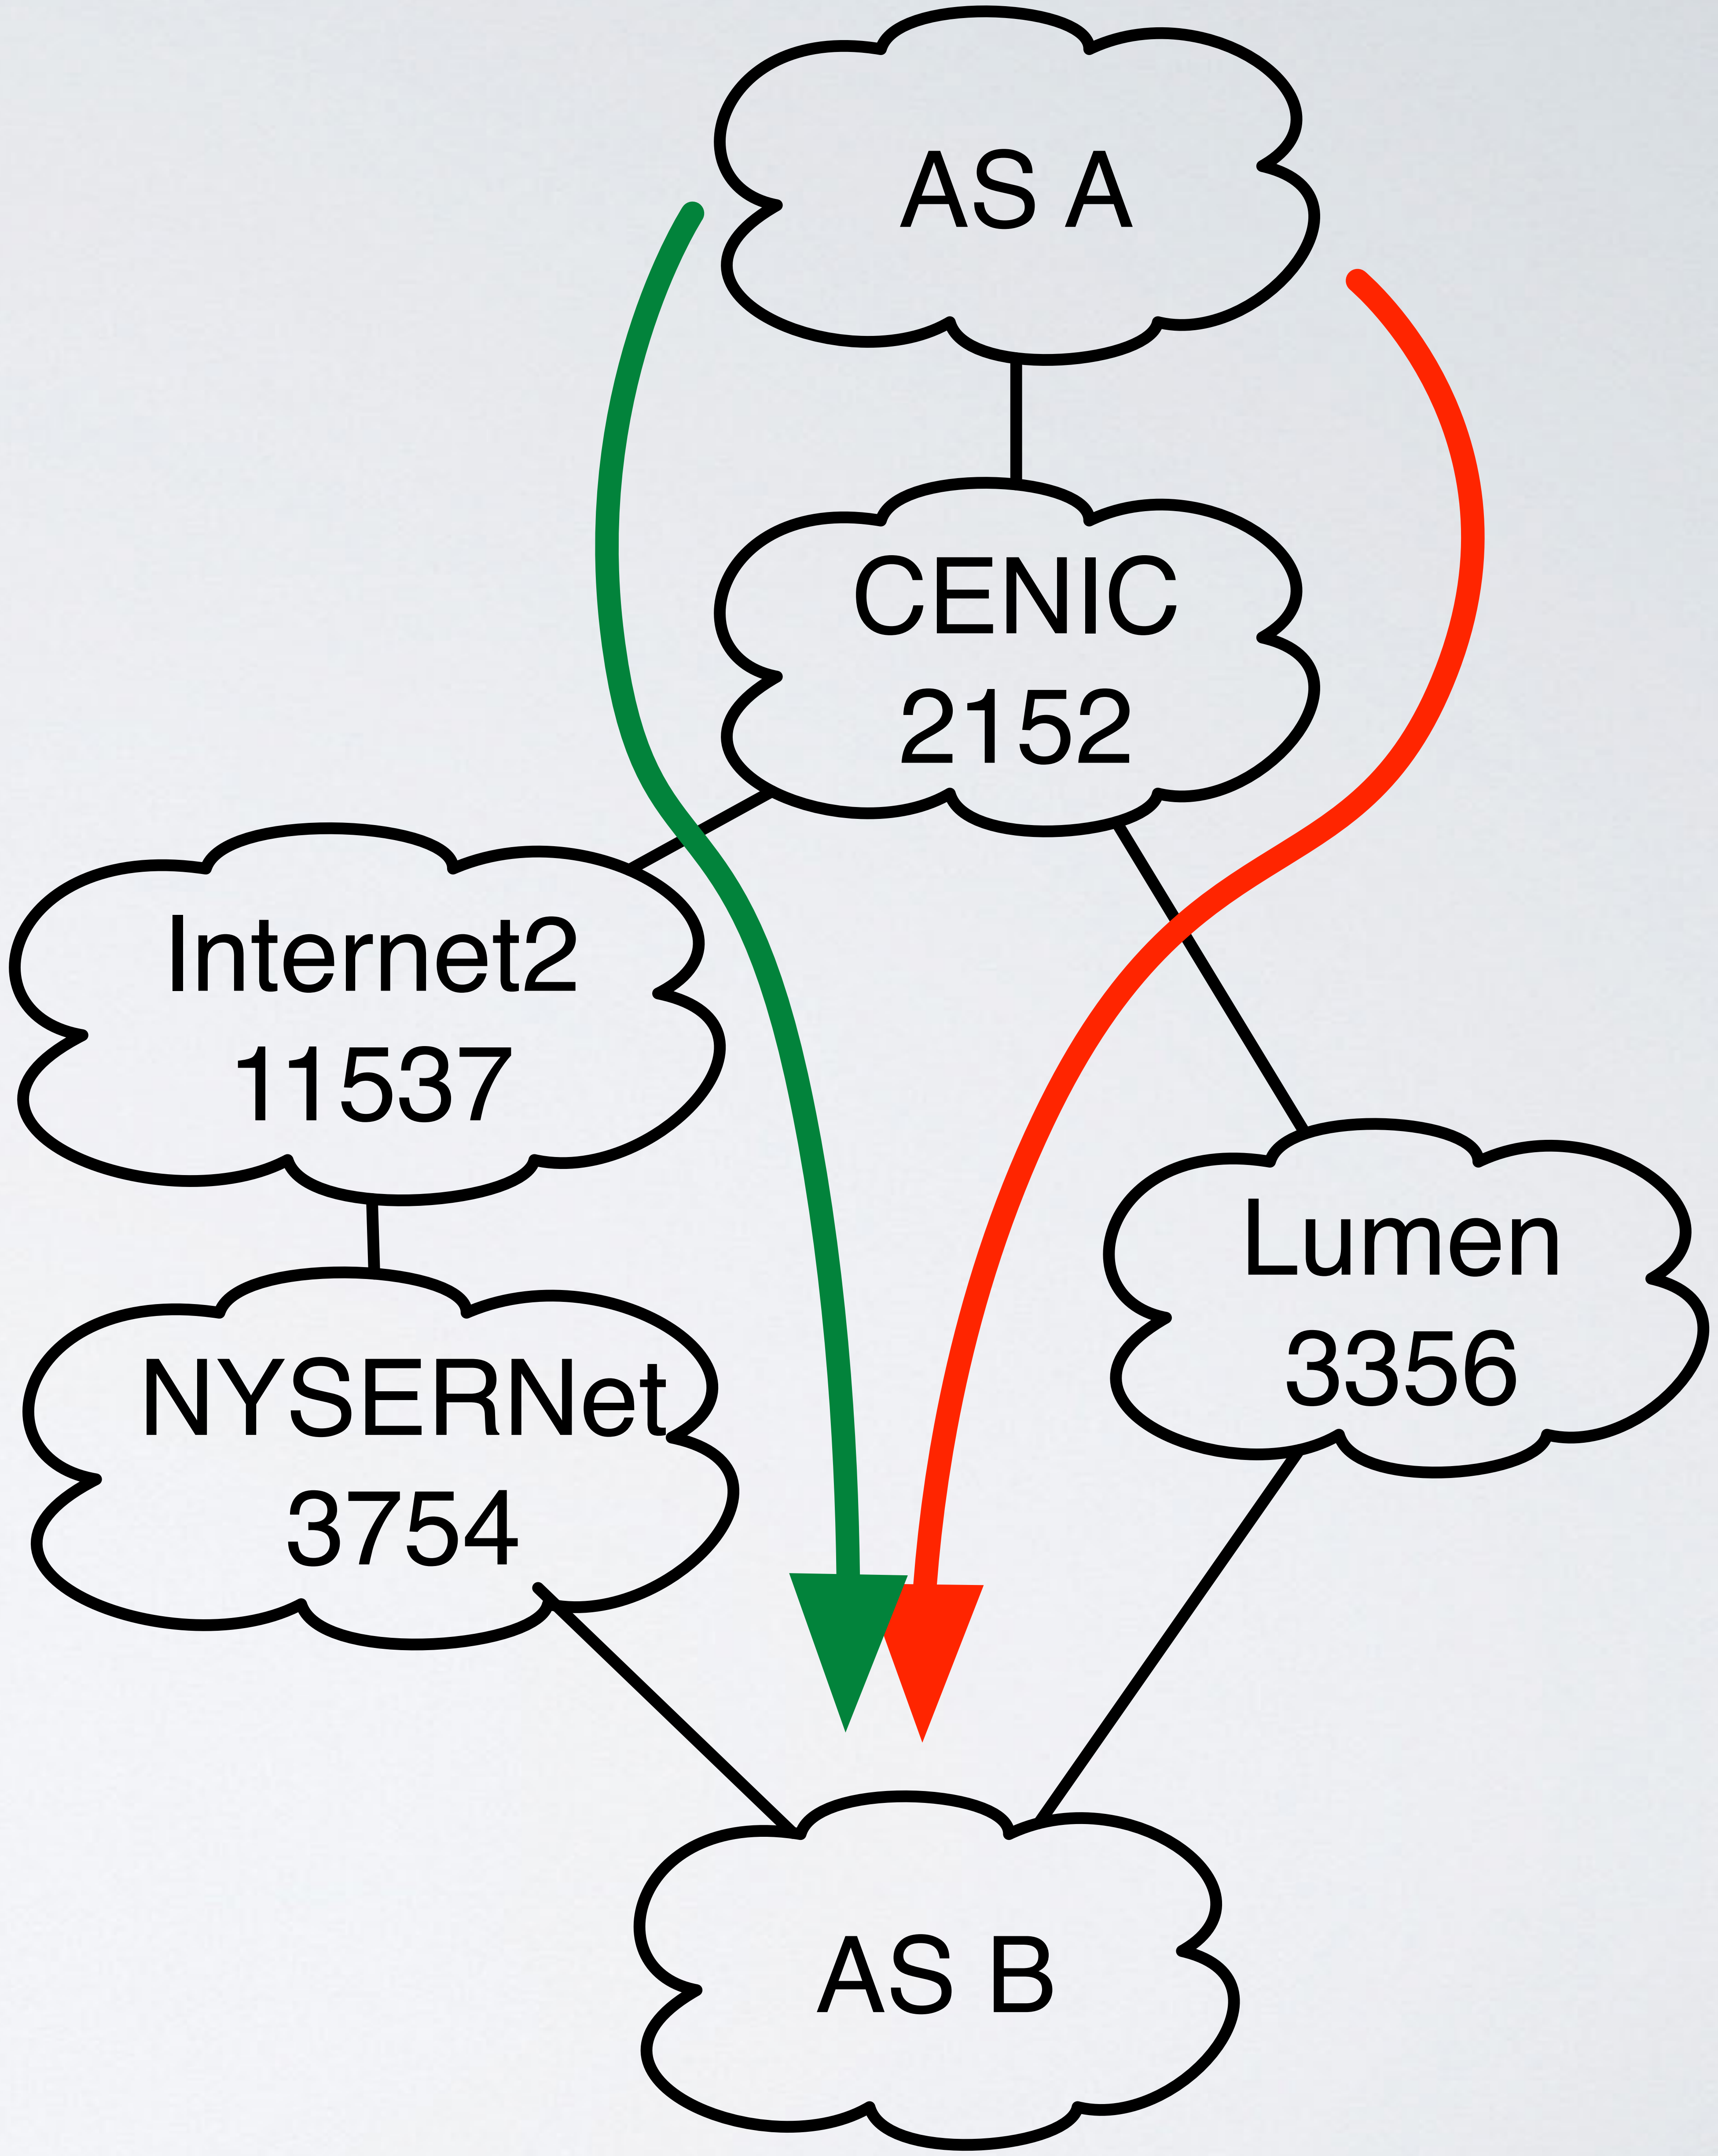

Challenge: Limited Visibility of Routing Policies

2,659 R&E ASes

RouteViews and RIPE RIS:
27 R&E ASes provide a public BGP view (1.0%)

RIPE Atlas:
215 R&E ASes provide an Atlas VP (8.1%)

Different systems in an AS can
experience different routing policies

Investigating R&E Routing Policies at Scale

Requirements

1. Ability to originate routes whose properties we'll measure

2. Measure global effects of route announcements without requiring measured ASes to do anything to assist the inference beyond responding to a ping

Approach

Internet2 Measurement System connected to two networks.

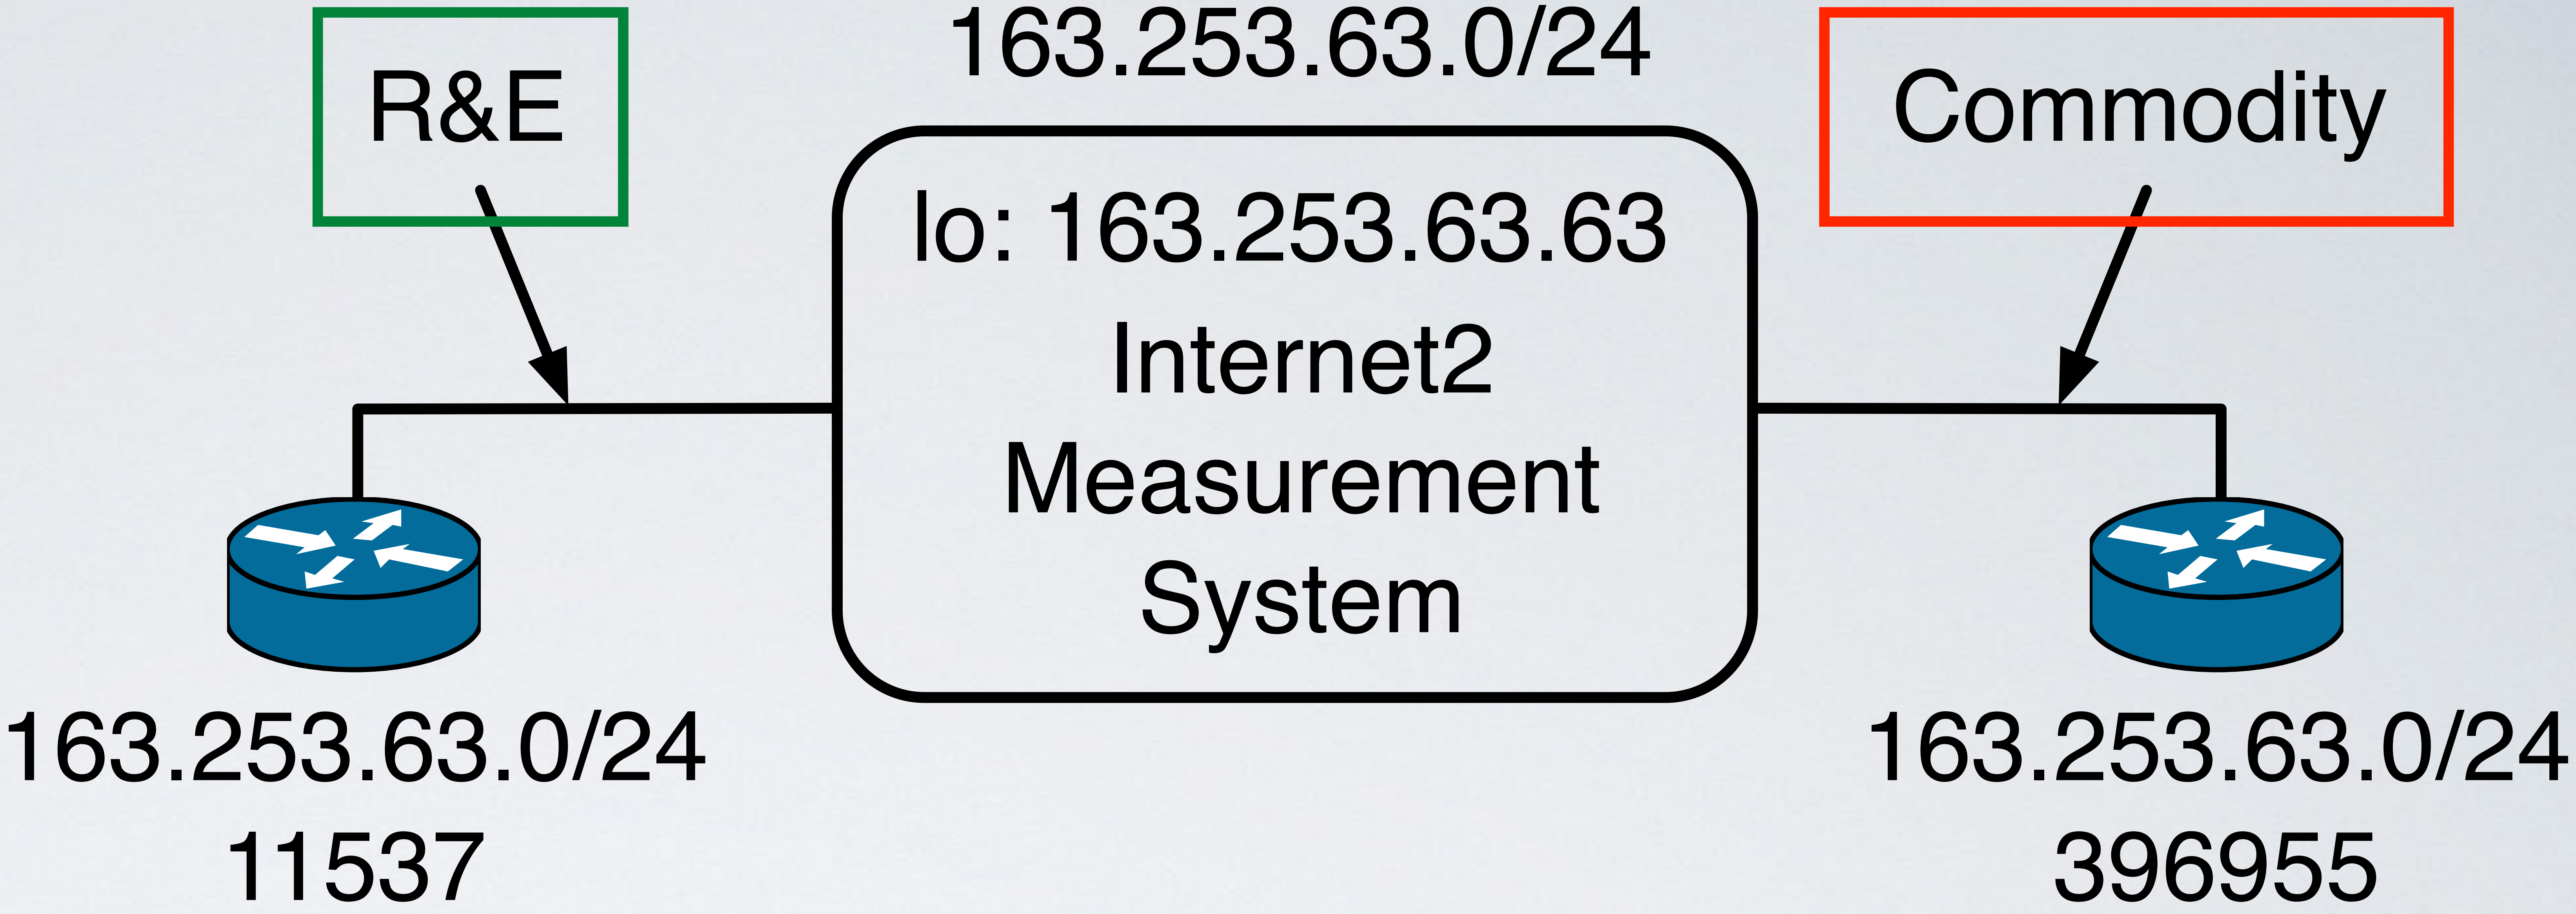

Approach

Internet2 Measurement System connected to two networks.

Both networks announce the same prefix, but with a different origin AS

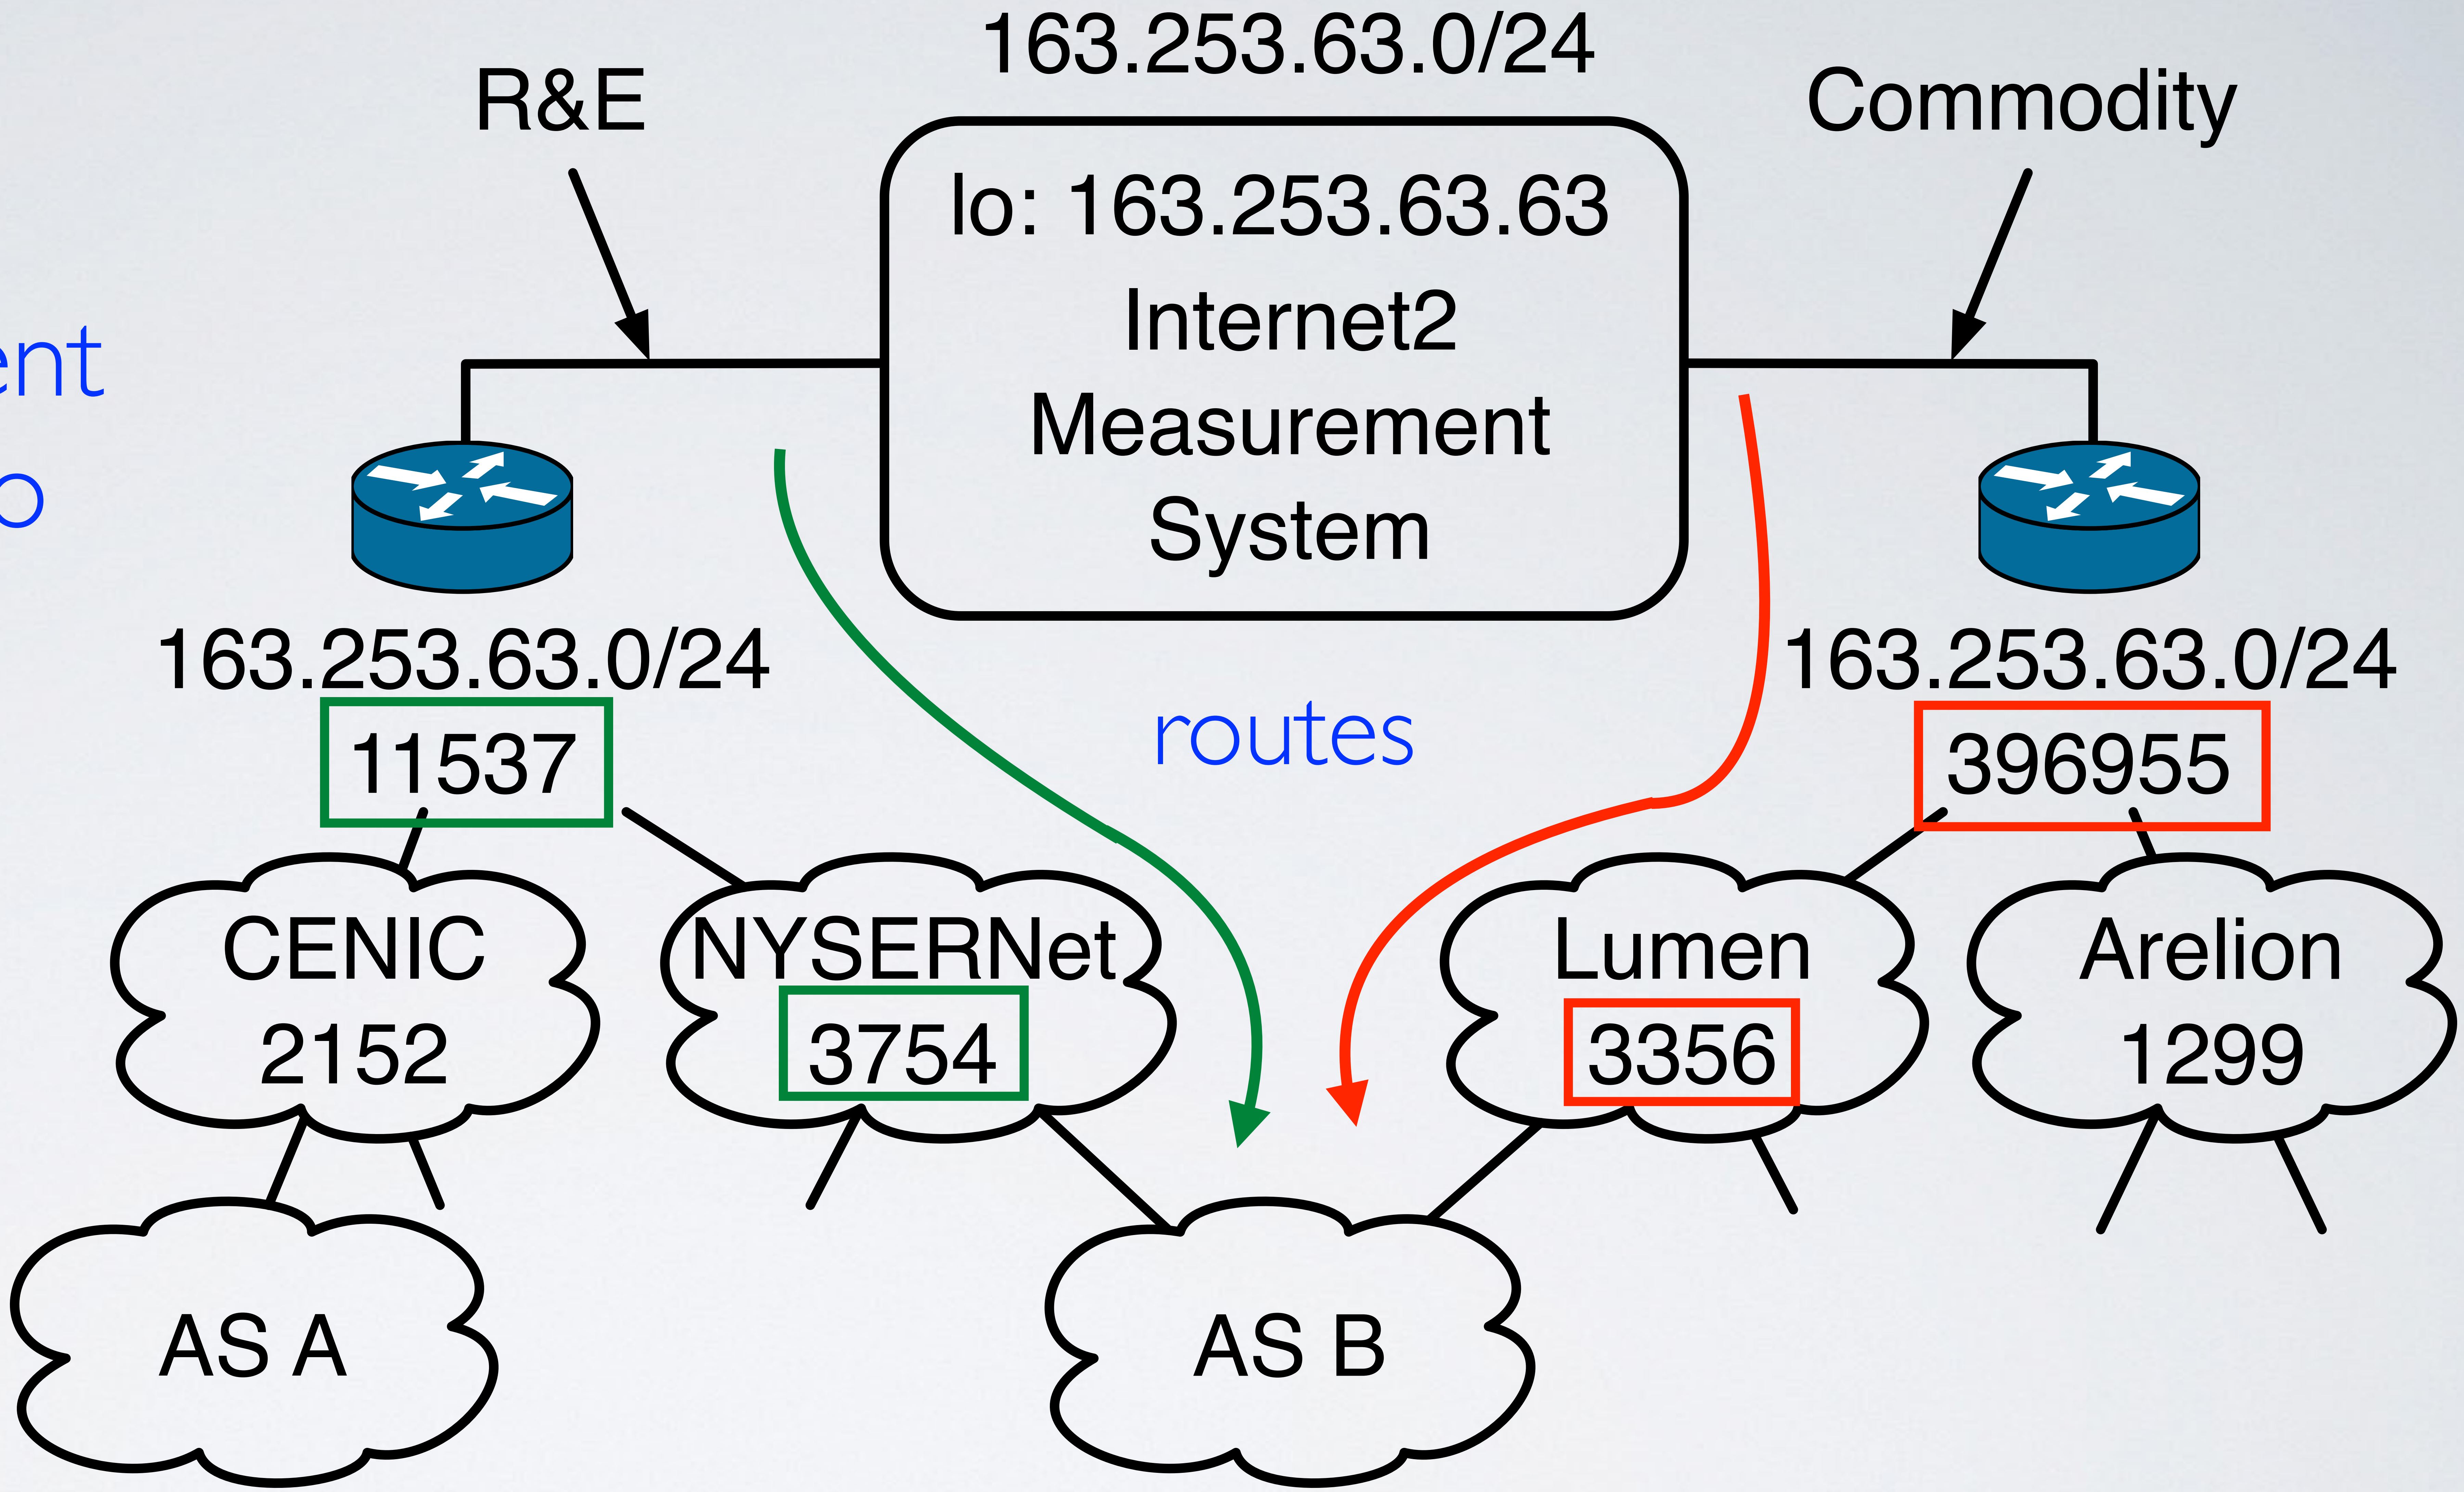

R&E route: 3754 11537 163.253.63.0/24

Commodity route: 3356 396955 163.253.63.0/24

Approach

Internet2 Measurement System connected to two networks.

Both networks announce the same prefix, but with a different origin AS

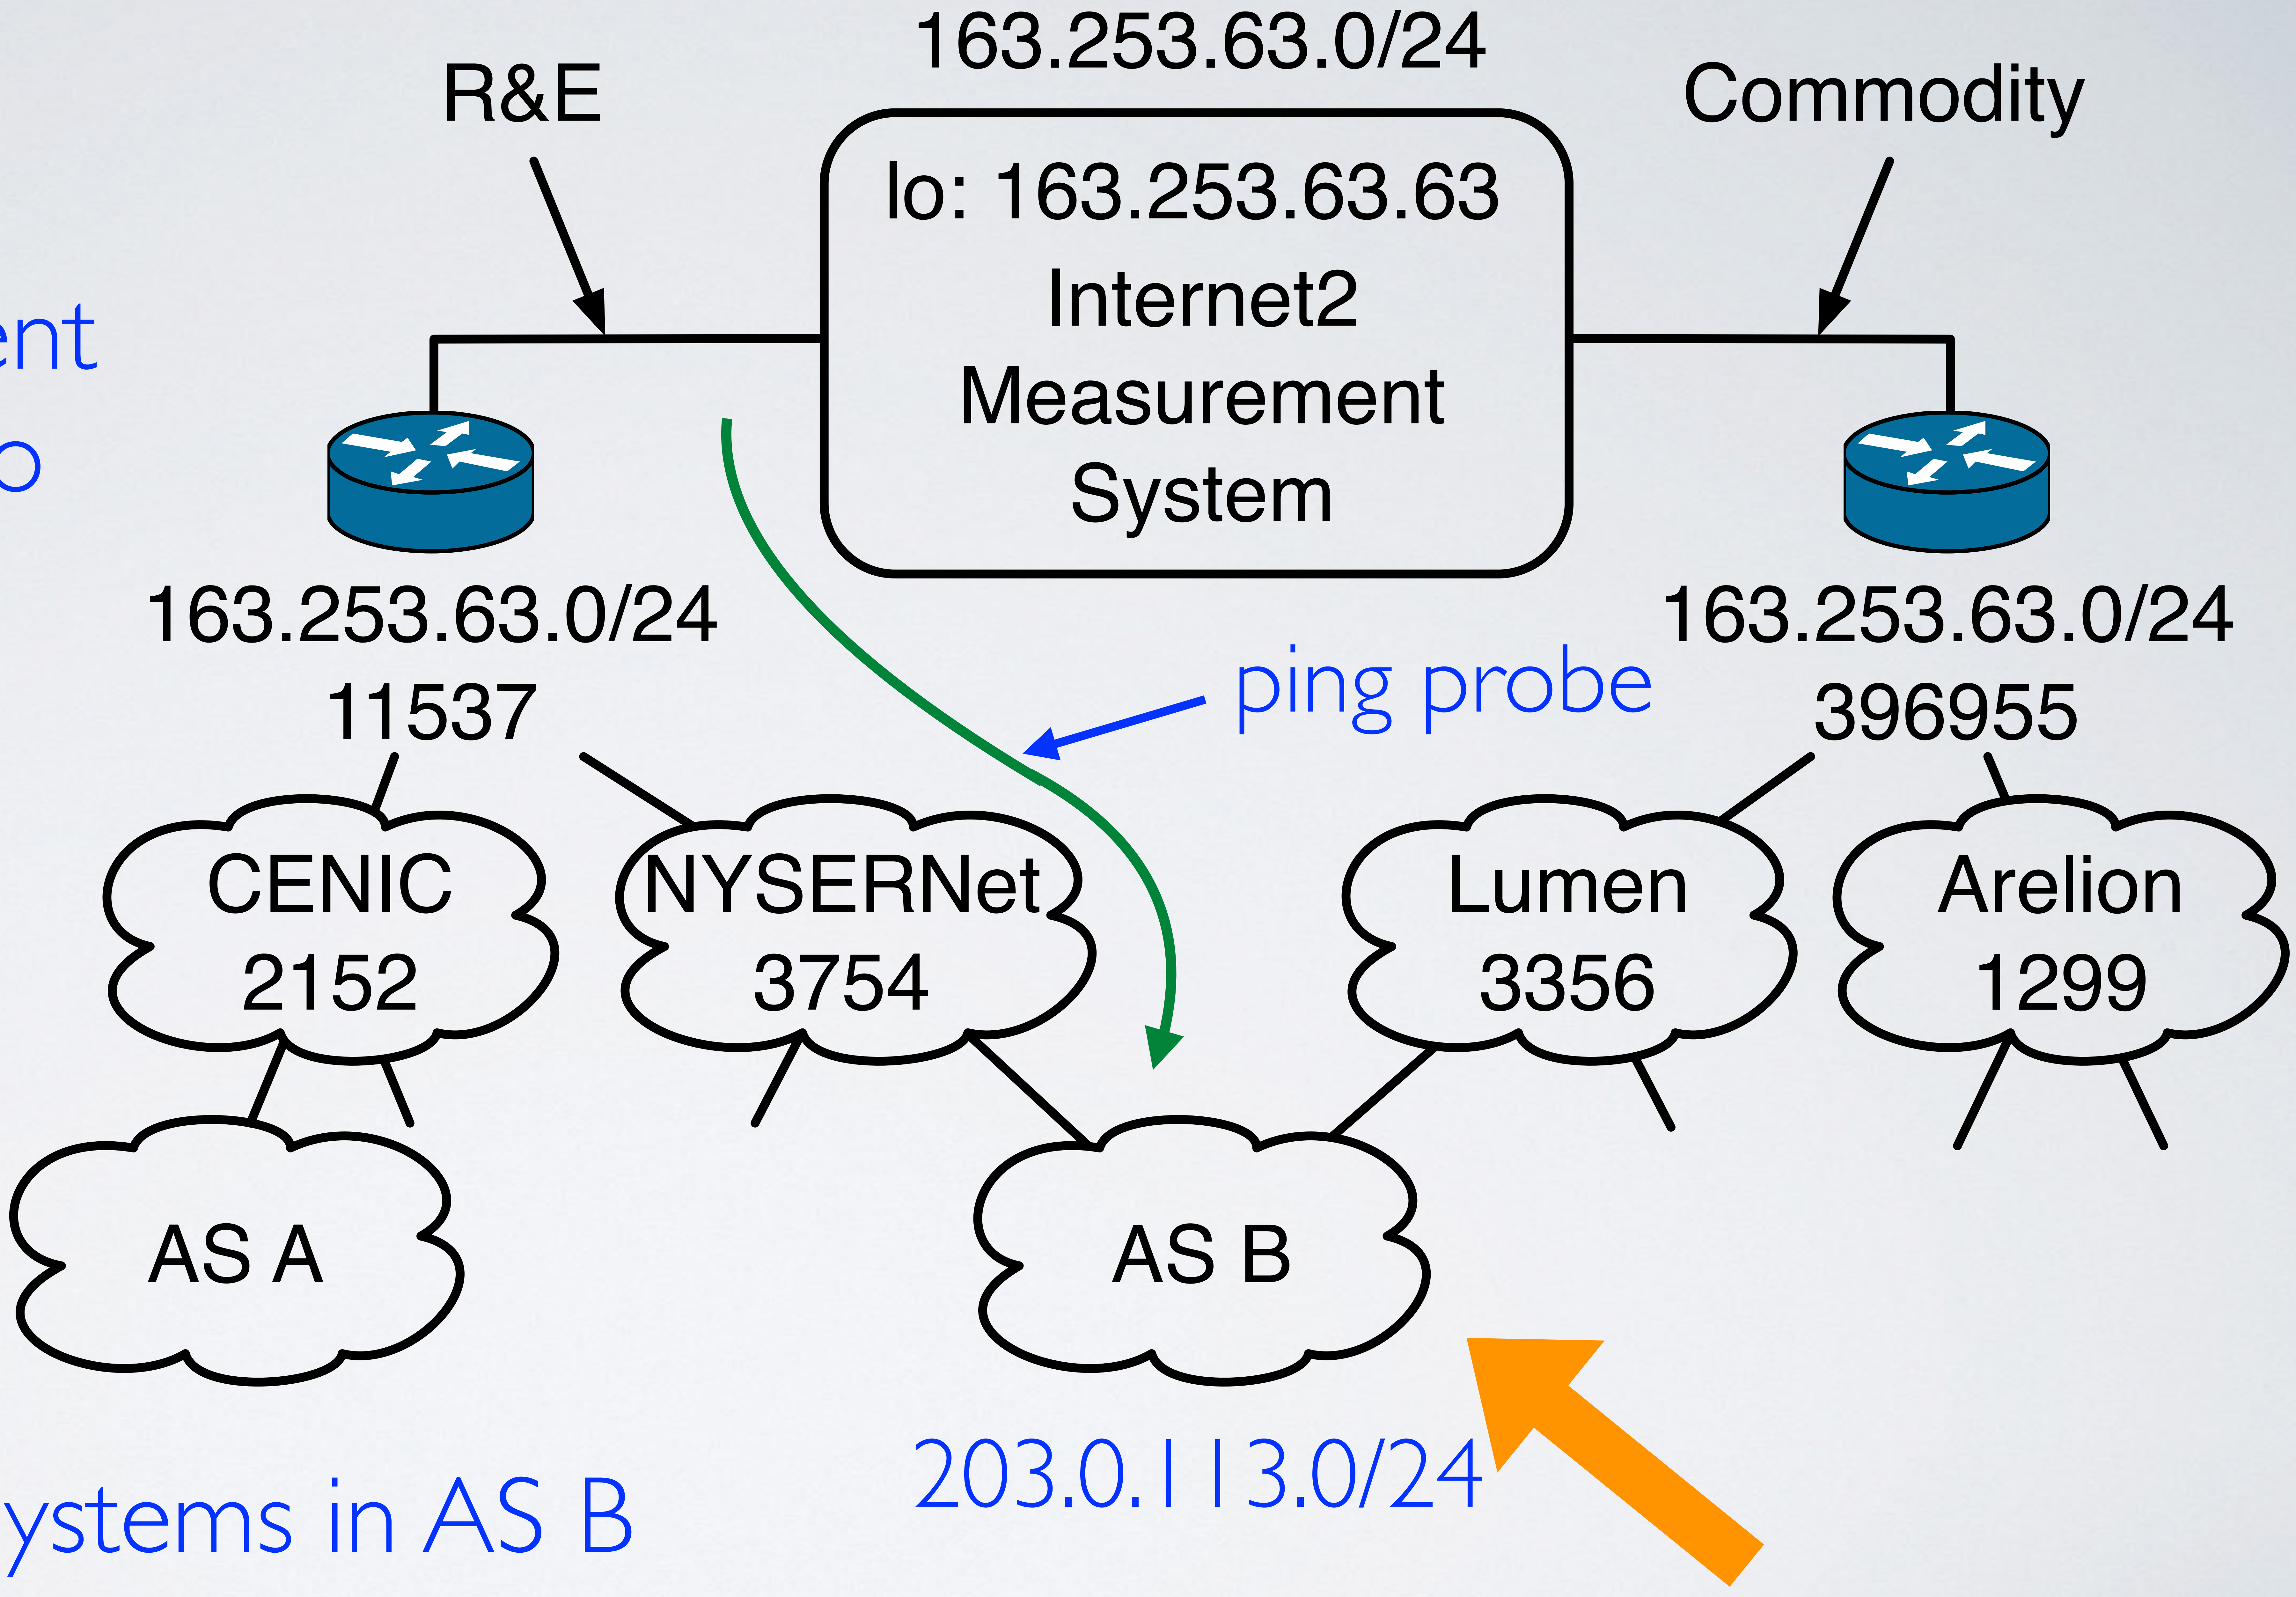

Ping-responsive systems in AS B can indicate B's routing policy

Approach

Internet2 Measurement System connected to two networks.

Both networks announce the same prefix, but with a different origin AS

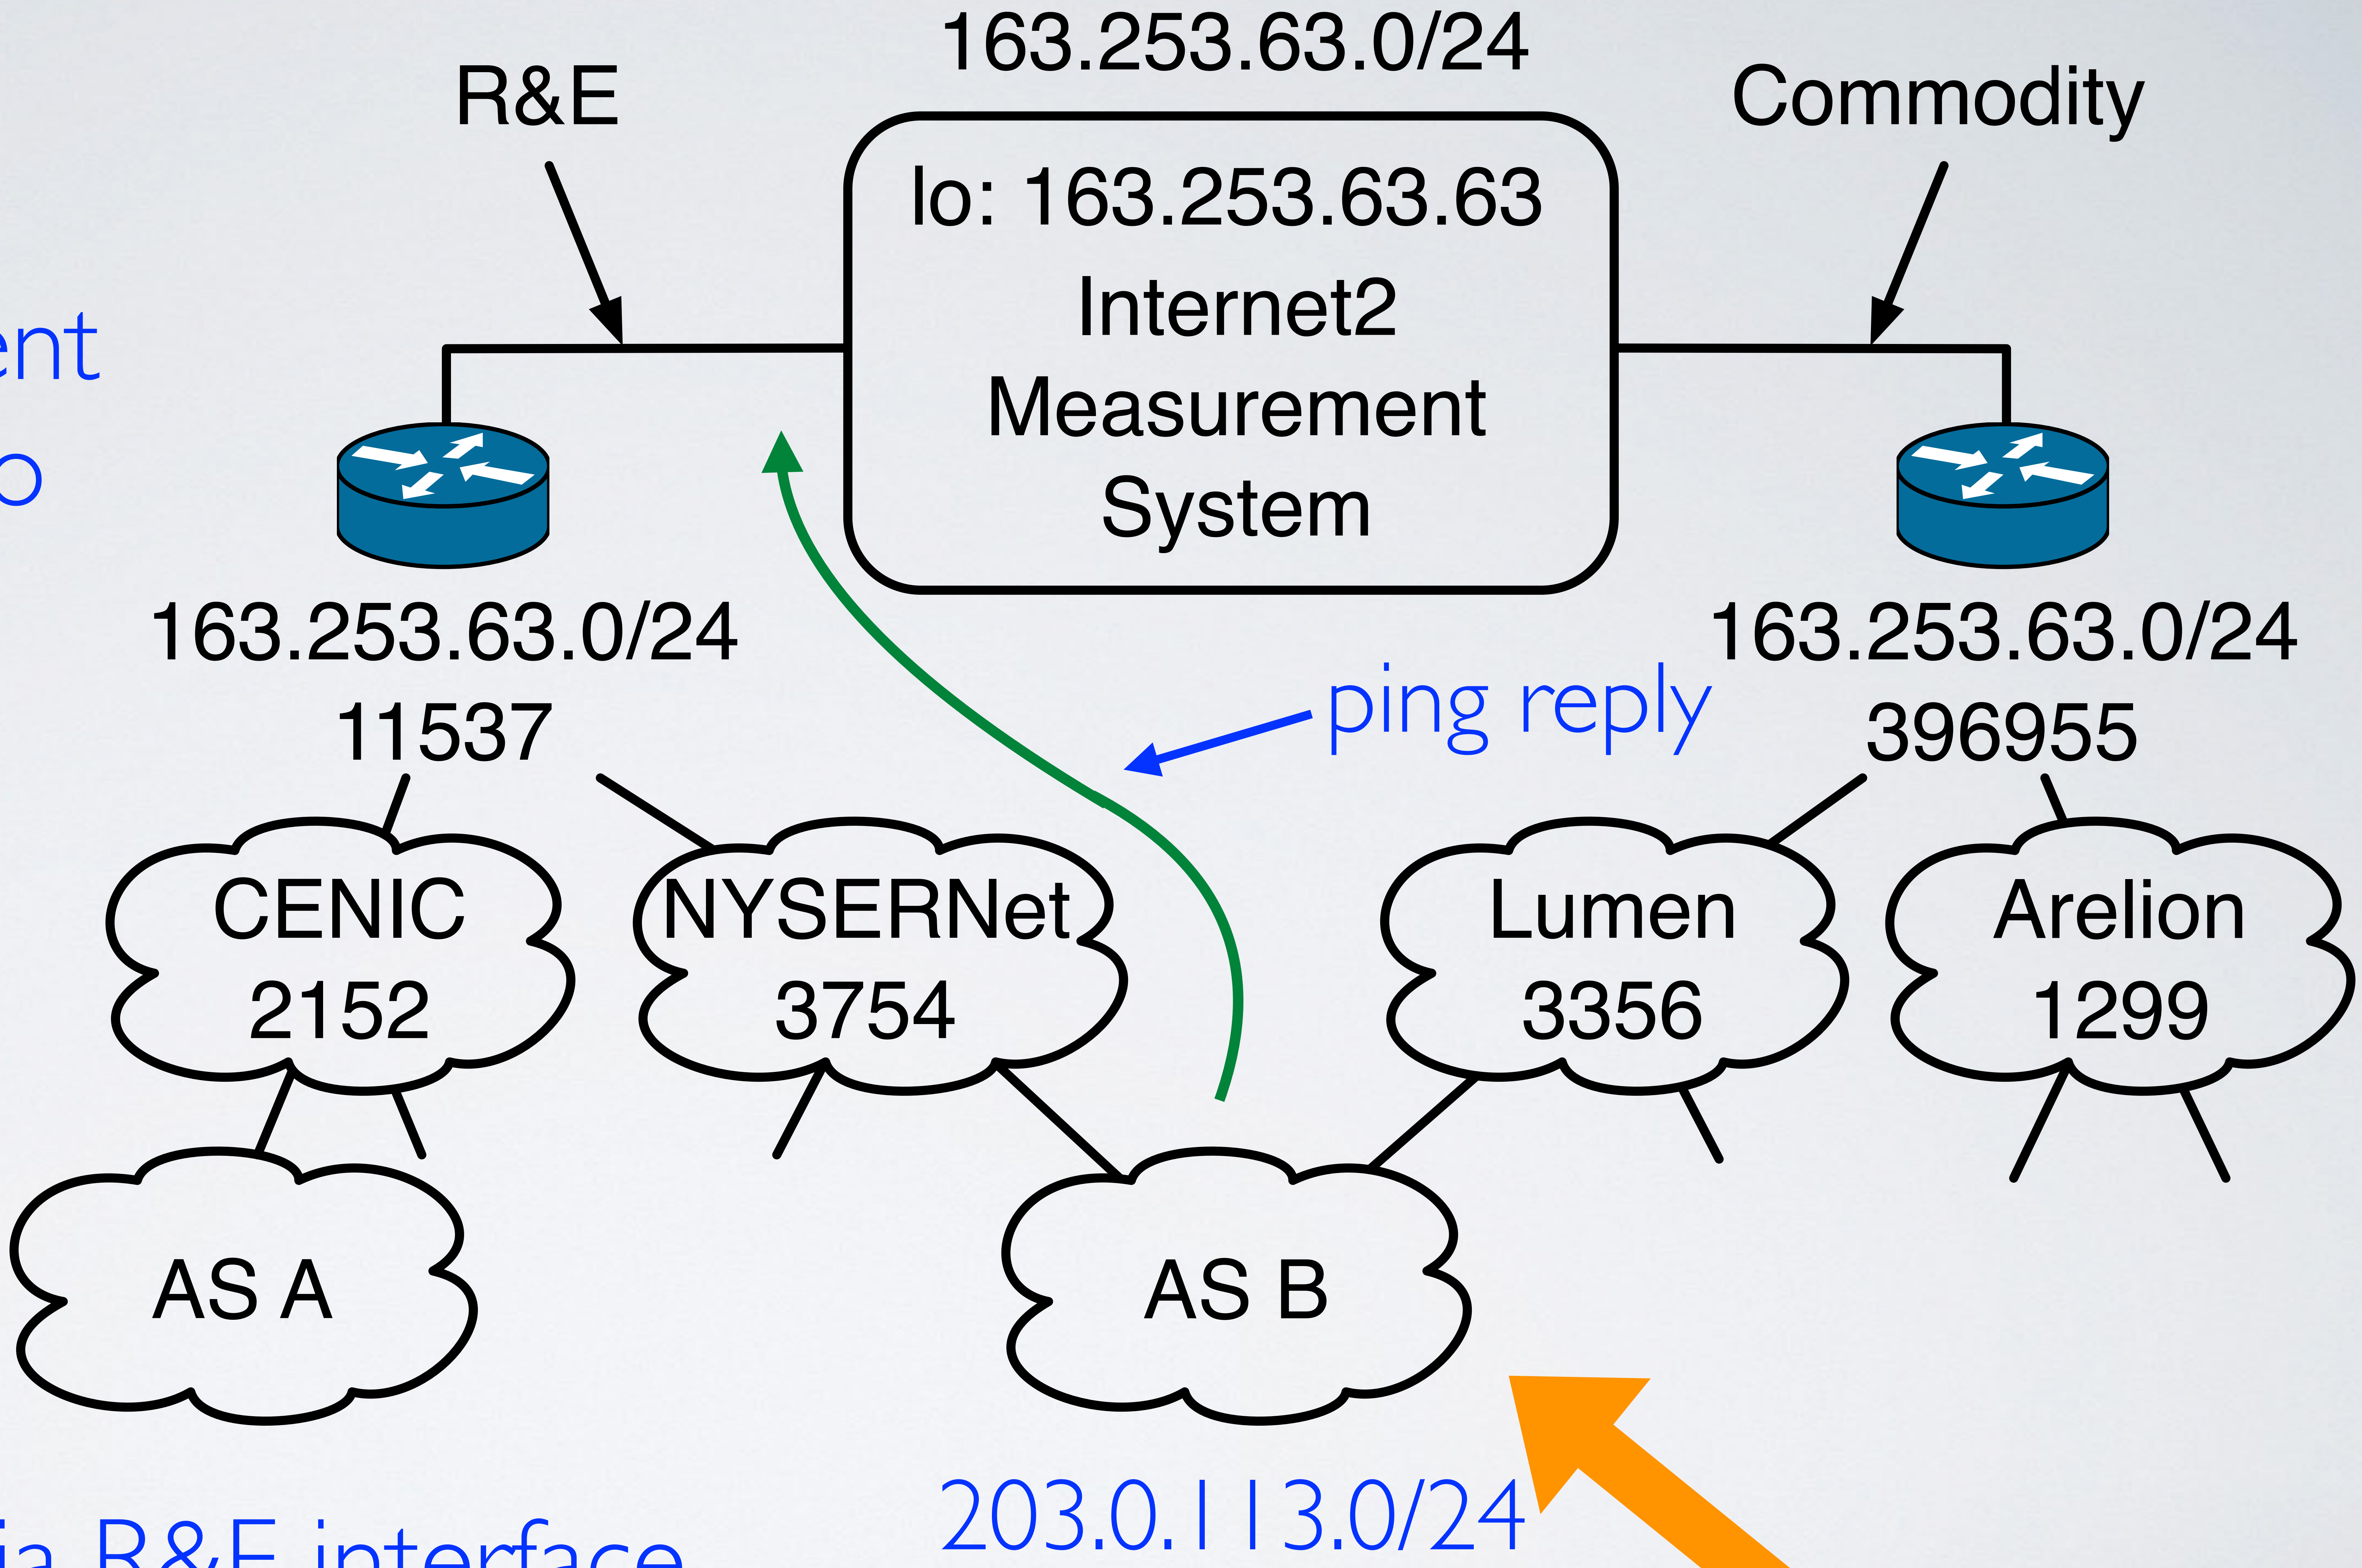

Replies received via R&E interface implies AS B chose an R&E route

Approach

Internet2 Measurement System connected to two networks.

Both networks announce the same prefix, but with a different origin AS

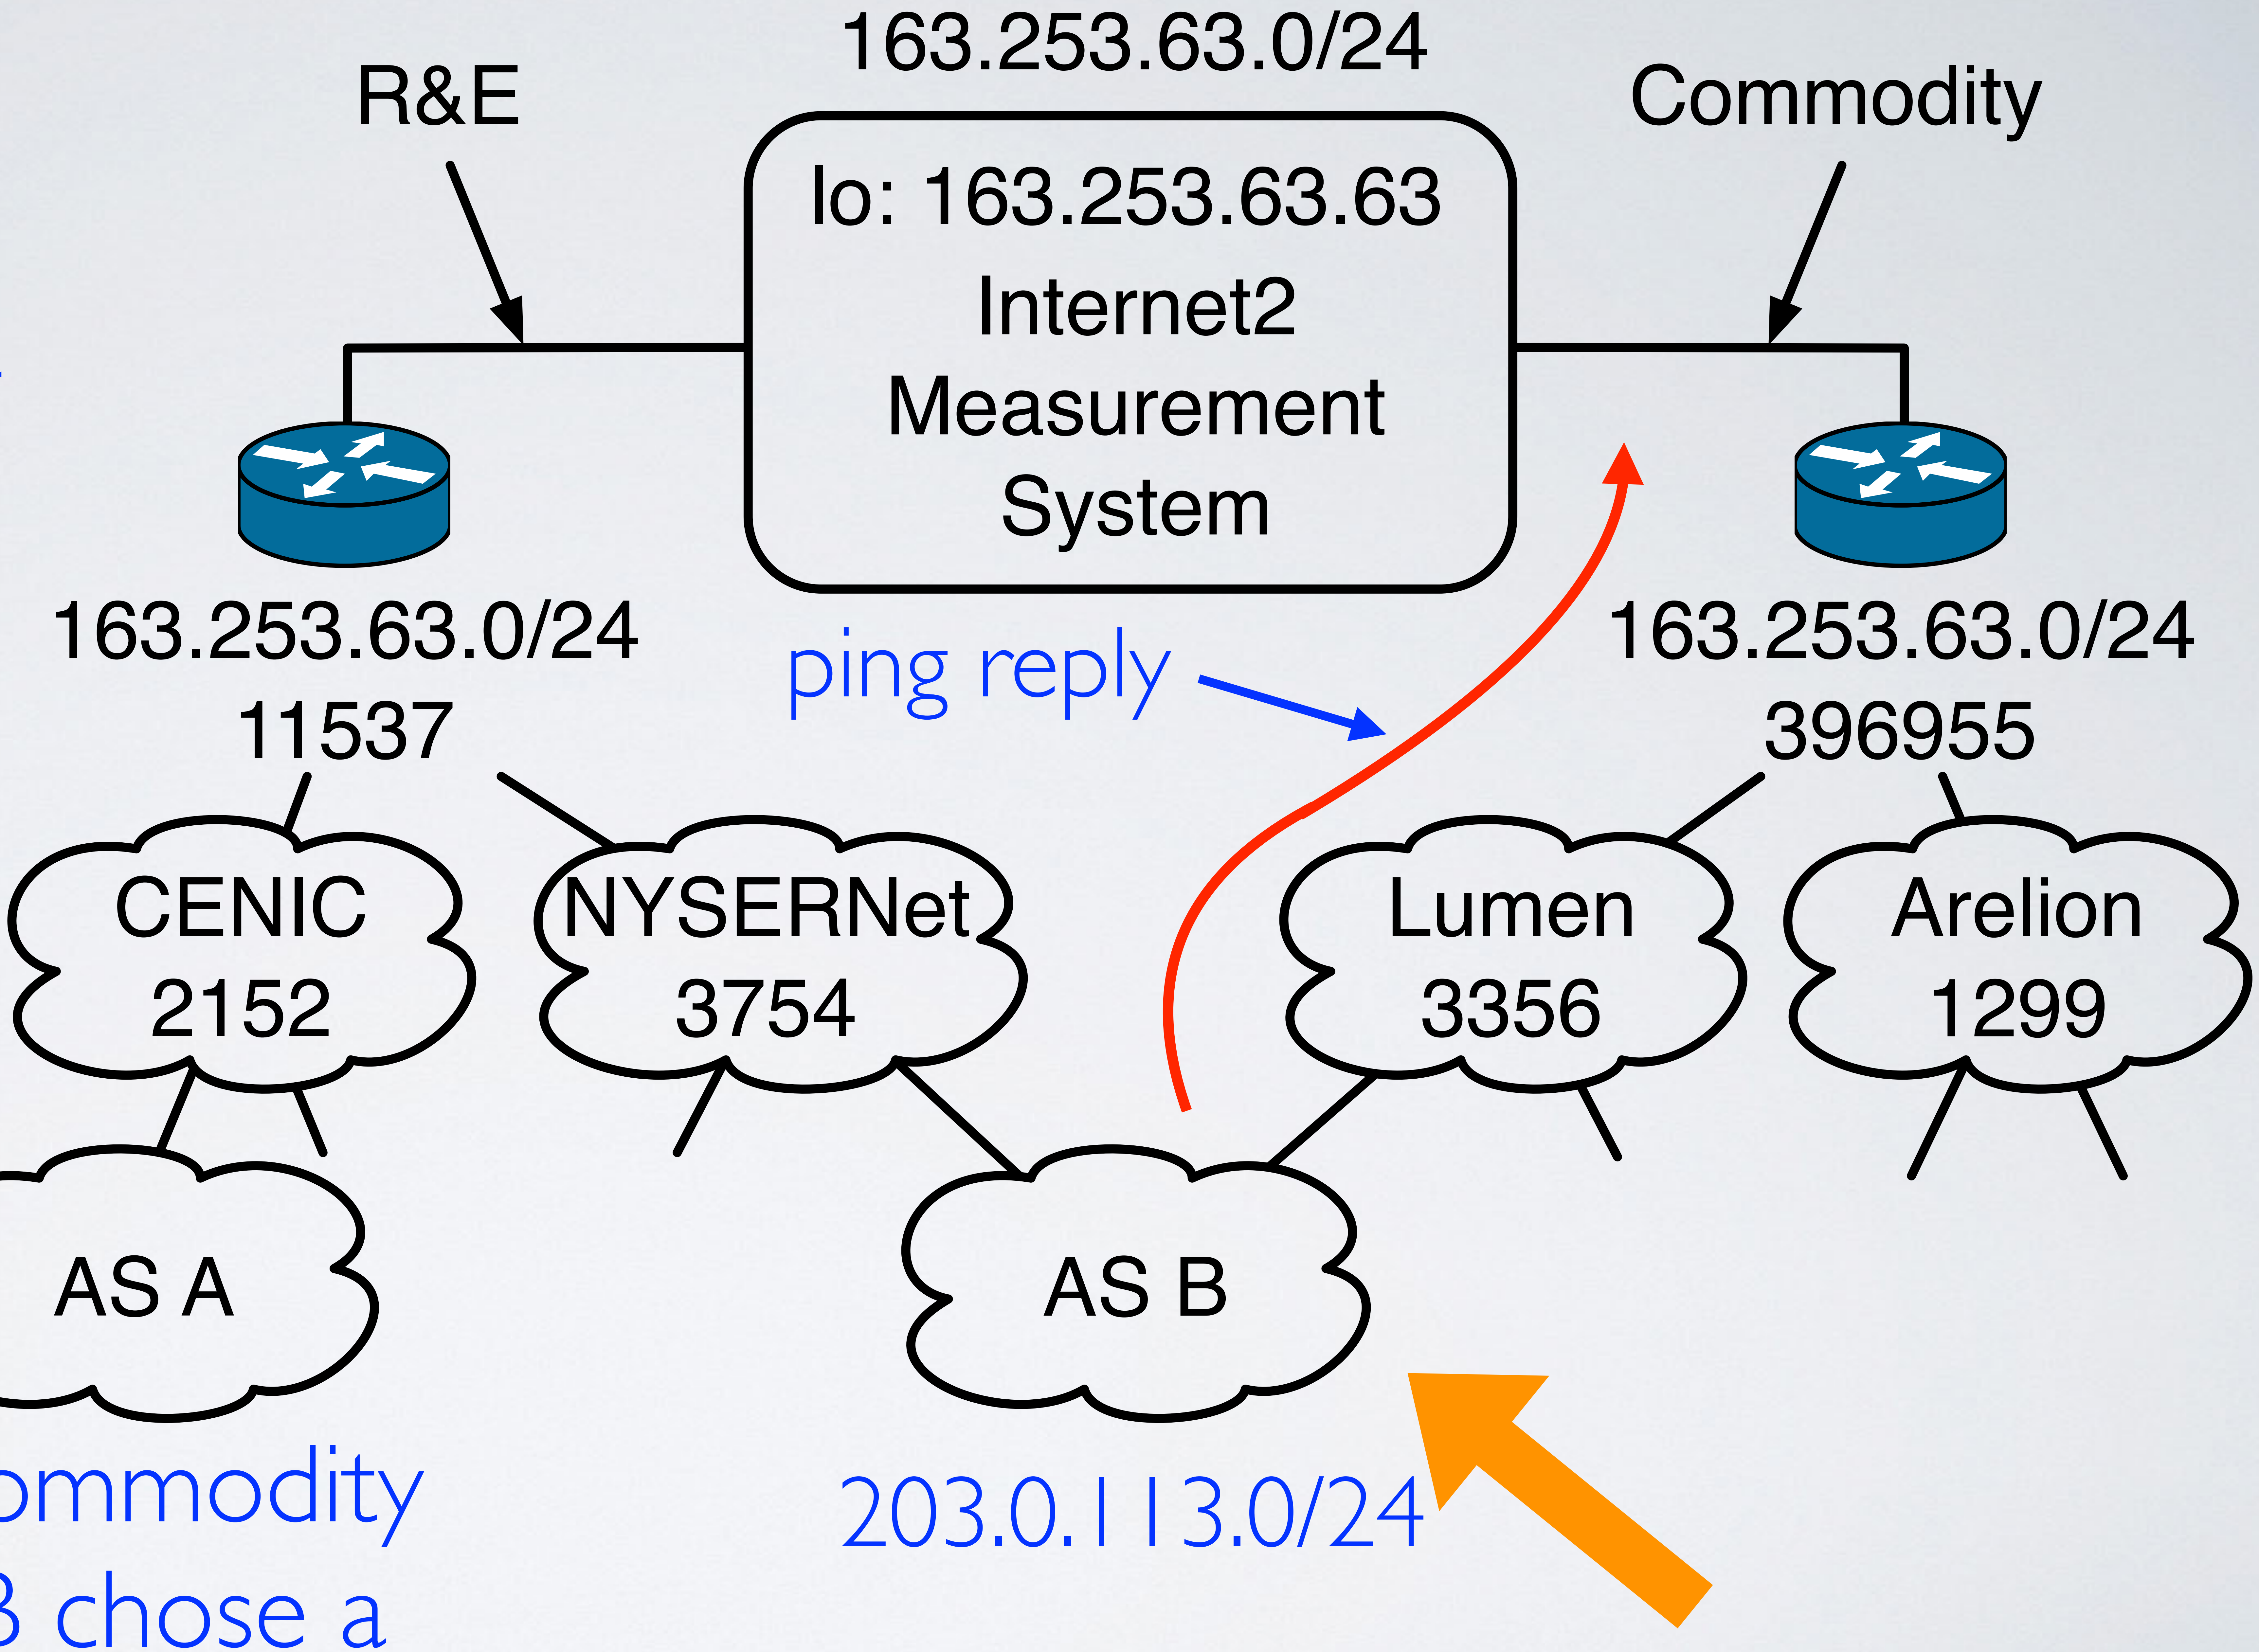

Replies received via commodity interface implies AS B chose a commodity route

Do R&E ASes assign R&E routes a higher localpref?

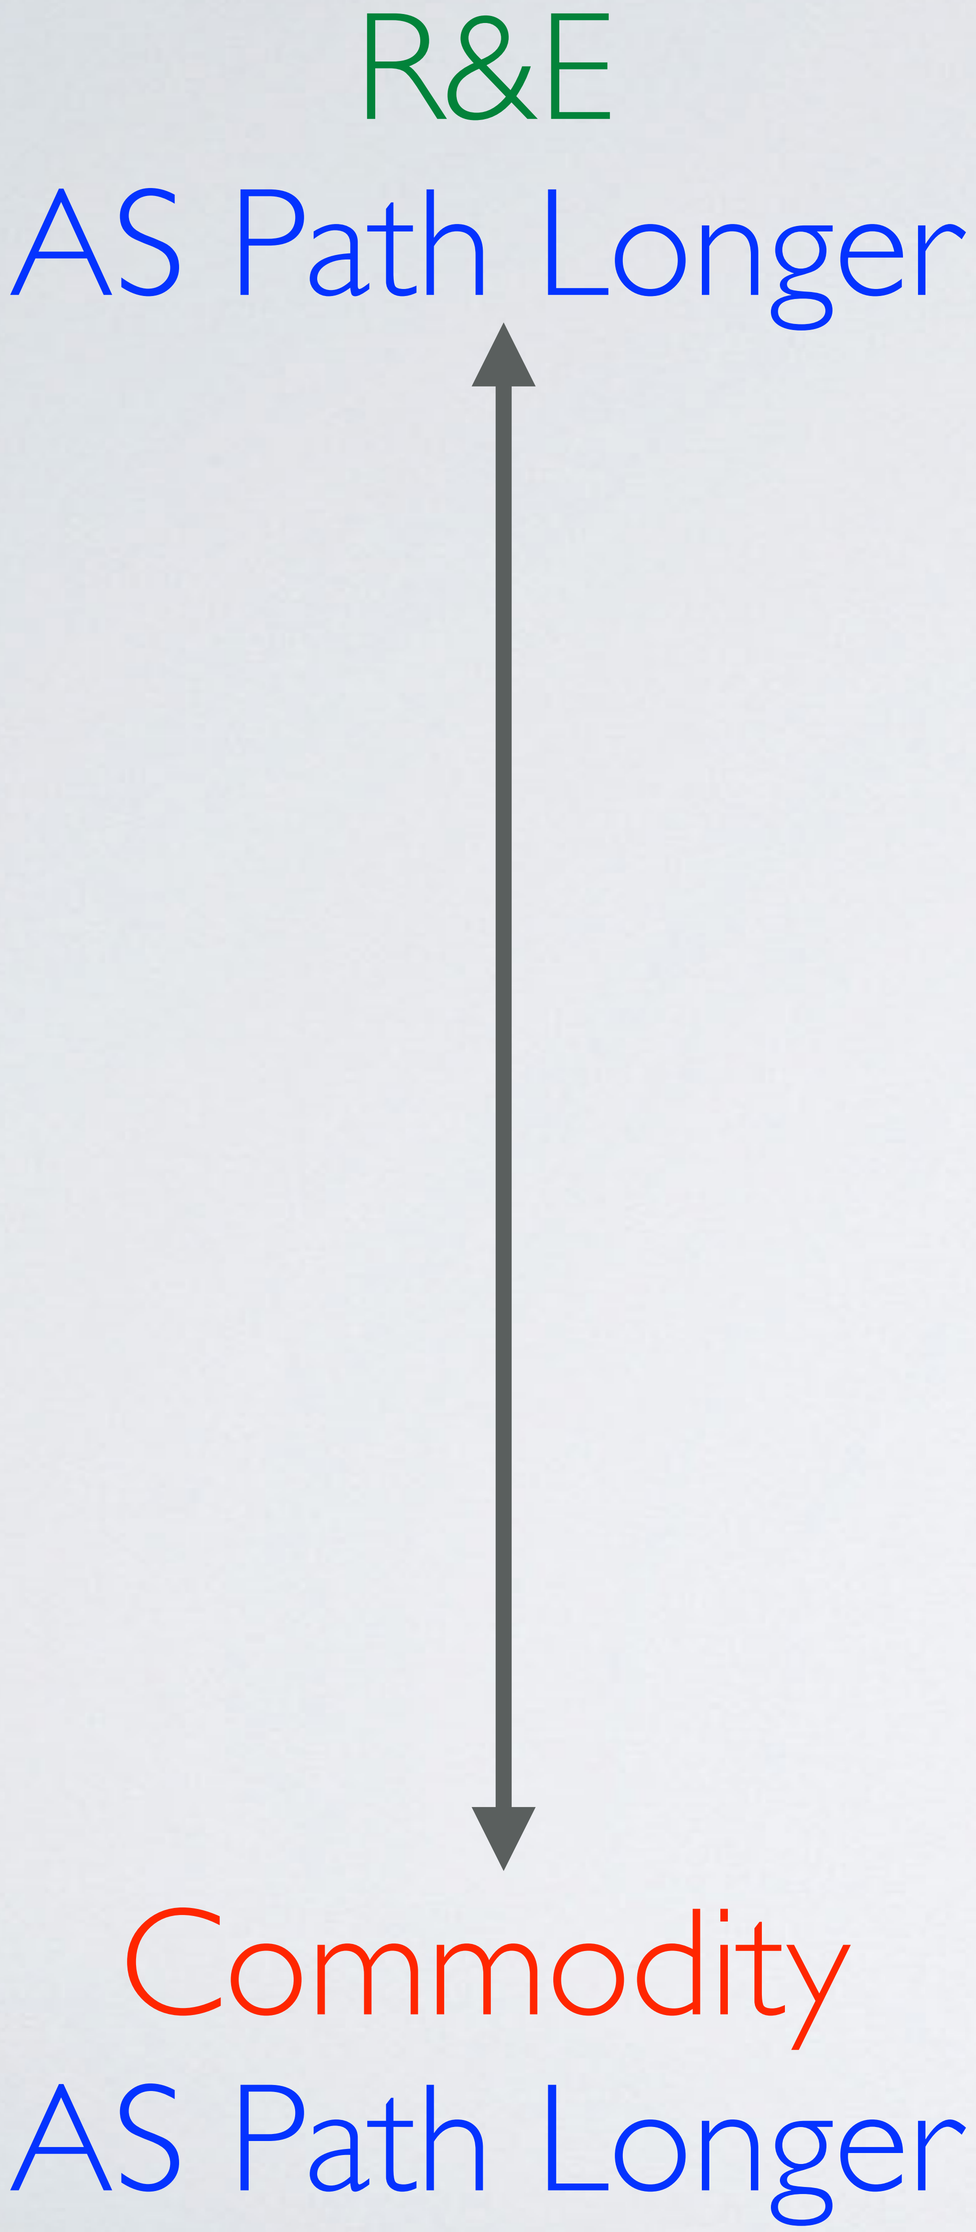

R&E
AS Path Longer

Commodity
AS Path Longer

- prepend R&E **4 times**
- prepend R&E **3 times**
- prepend R&E **2 times**
- prepend R&E **once**
- no prepending of either route
- prepend commodity **once**
- prepend commodity **2 times**
- prepend commodity **3 times**
- prepend commodity **4 times**

Routers that localpref R&E will select R&E route even when AS path is long

Routers tie-breaking with AS path length will switch from commodity to R&E

Routers that localpref commodity will select commodity route even when AS path is long

Result Summary - June 5th 2025

Inference	Prefixes		ASes	
Always R&E	9,758	**80.8%**	1,940	**75.3%**
Always commodity	840	7.0%	353	13.7%
Switch to R&E	1,103	9.1%	322	12.5%
Switch to commodity	3	0.0%	3	0.1%
Mixed R&E and commodity	371	3.1%	228	8.8%
Oscillating	2	0.0%	2	0.1%
Total:	12,077		2,578	

Majority always used the R&E route
(regardless of prepend configuration)

Result Summary - June 5th 2025

Inference	Prefixes		ASes	
Always R&E	9,758	80.8%	1,940	75.3%
Always commodity	840	**7.0%**	353	**13.7%**
Switch to R&E	1,103	9.1%	322	12.5%
Switch to commodity	3	0.0%	3	0.1%
Mixed R&E and commodity	371	3.1%	228	8.8%
Oscillating	2	0.0%	2	0.1%
Total:	12,077		2,578	

Small fraction always used the commodity route (regardless of prepend configuration)

Result Summary - June 5th 2025

Inference	Prefixes		ASes	
Always R&E	9,758	80.8%	1,940	75.3%
Always commodity	840	7.0%	353	13.7%
Switch to R&E	1,103	**9.1%**	322	**12.5%**
Switch to commodity	3	0.0%	3	0.1%
Mixed R&E and commodity	371	3.1%	228	8.8%
Oscillating	2	0.0%	2	0.1%
Total:	12,077		2,578	

Small fraction used route with shortest AS path
(assigned same localpref to both commodity and R&E routes)

See the paper for more details…

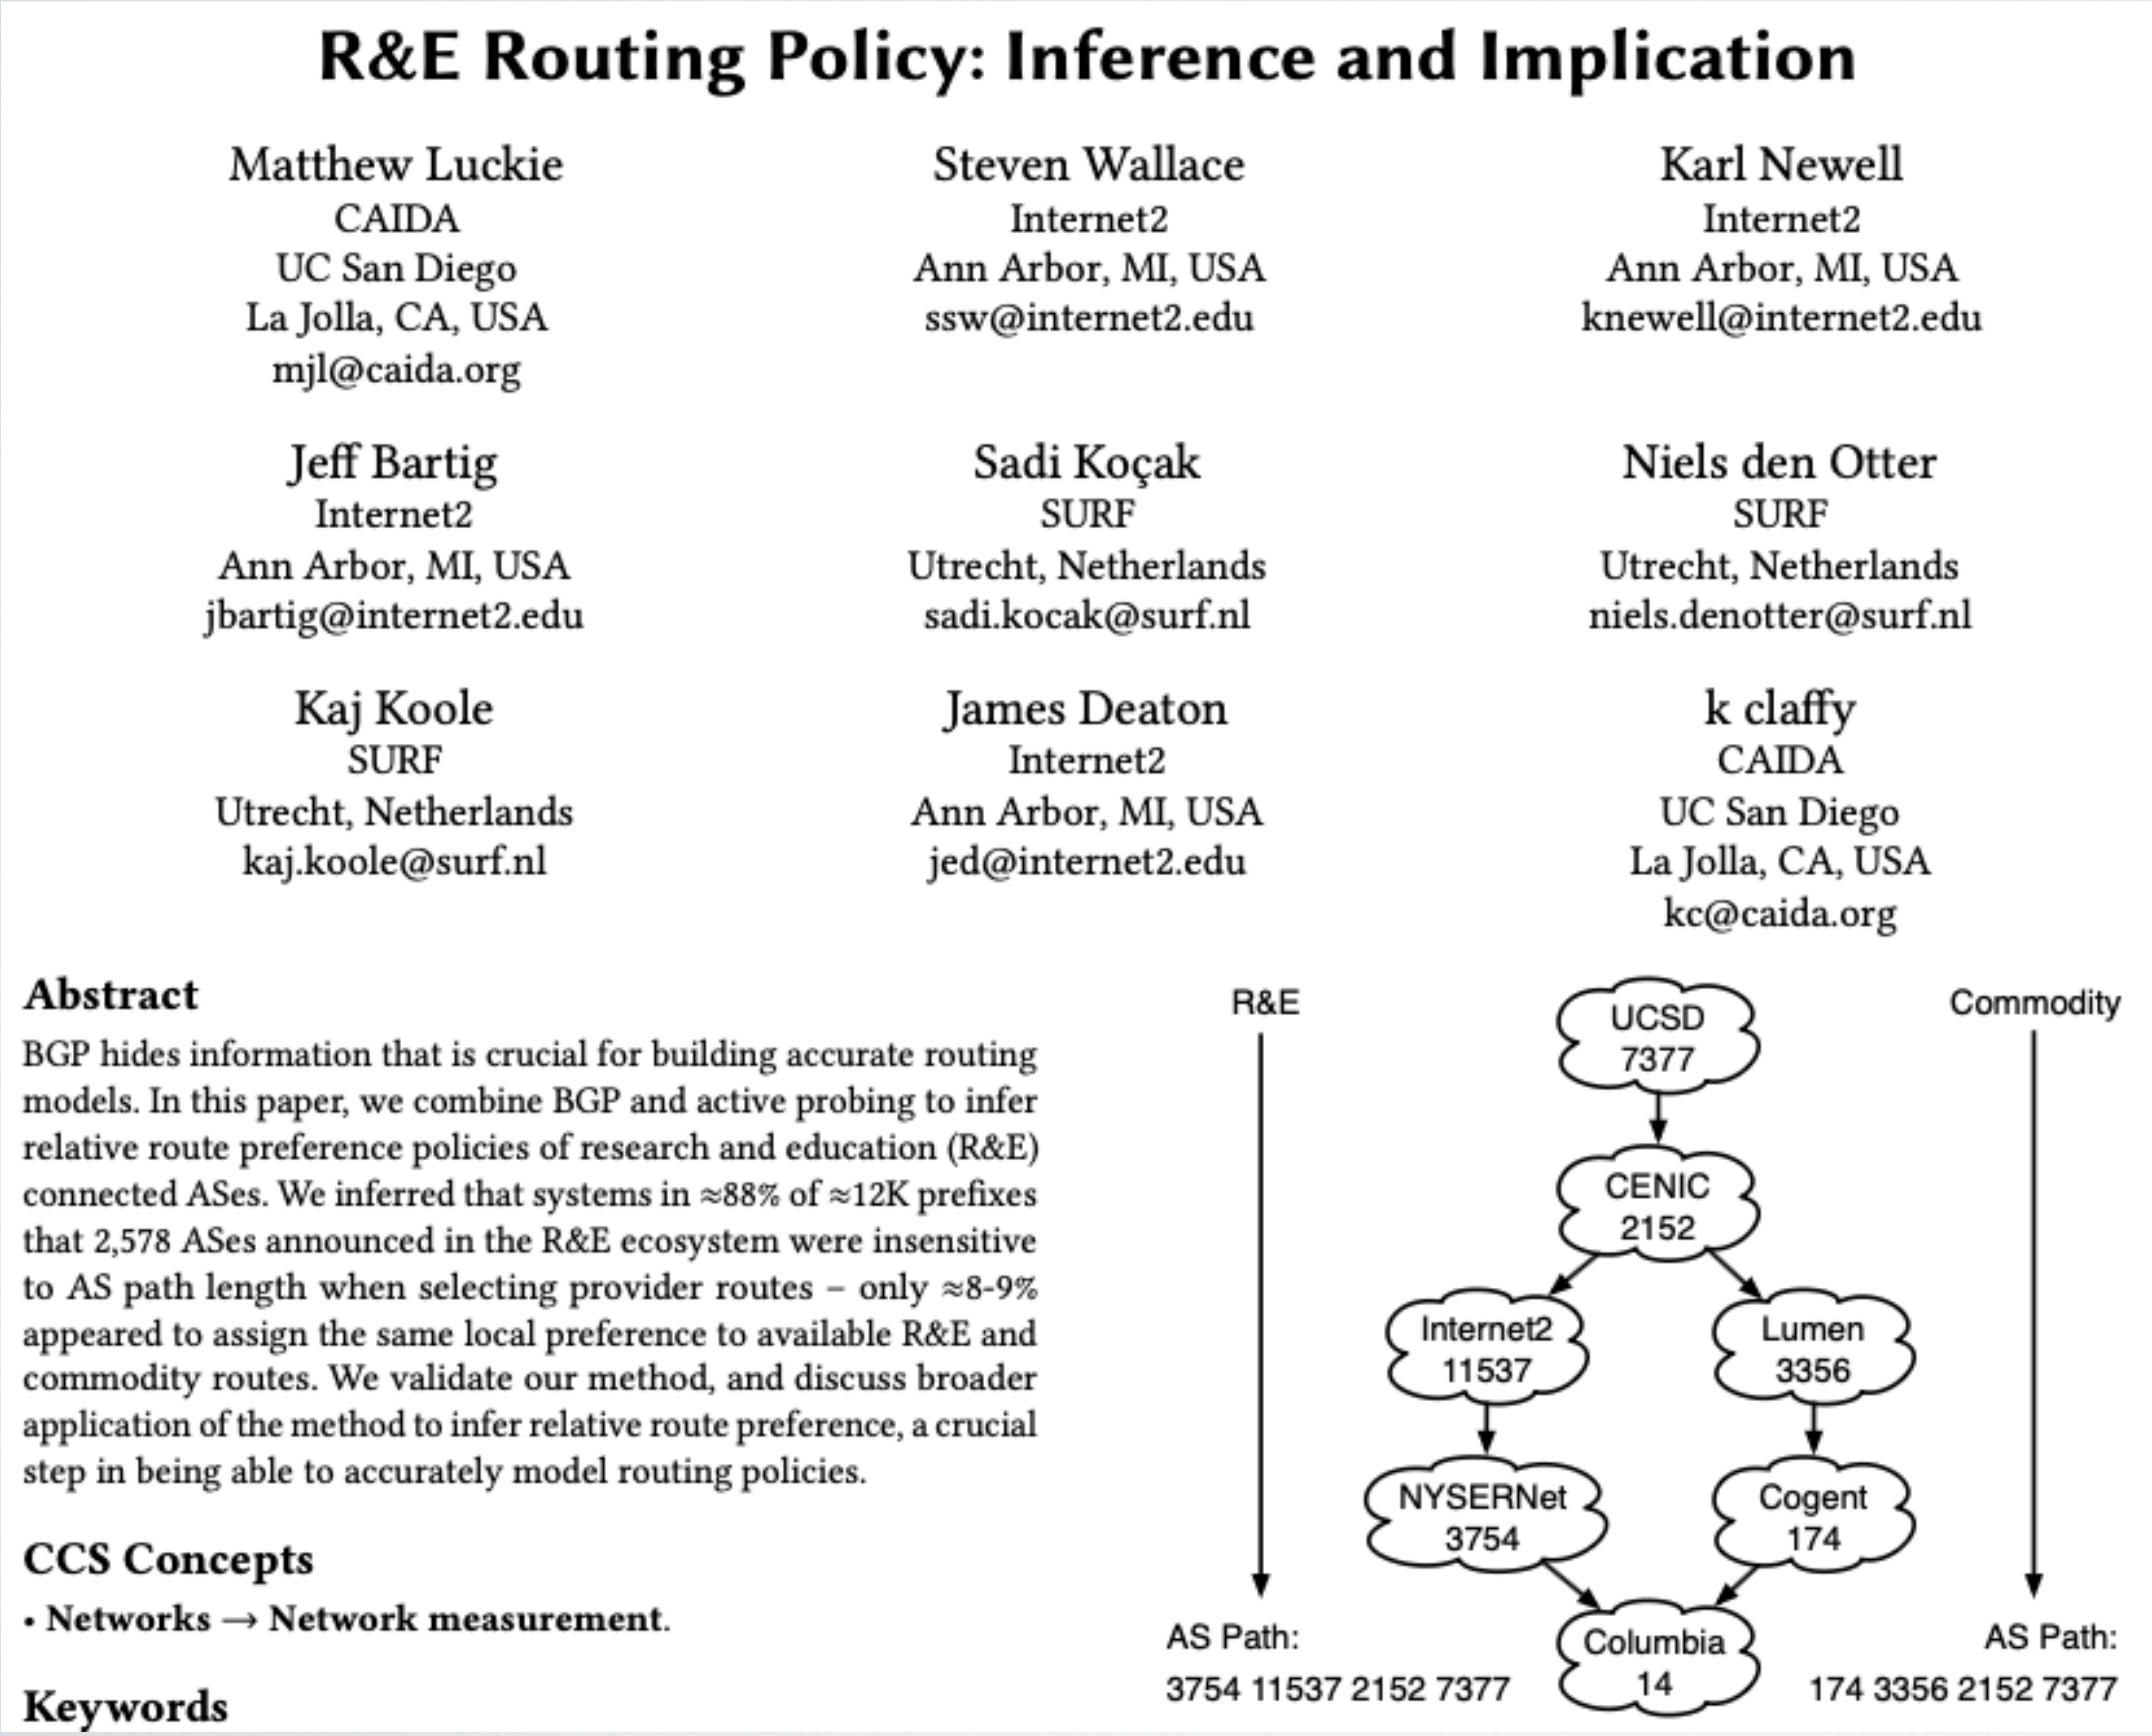

R&E Routing Policy: Inference and Implication

Matthew Luckie
CAIDA
UC San Diego
La Jolla, CA, USA
mjl@caida.org

Steven Wallace
Internet2
Ann Arbor, MI, USA
ssw@internet2.edu

Karl Newell
Internet2
Ann Arbor, MI, USA
knewell@internet2.edu

Jeff Bartig
Internet2
Ann Arbor, MI, USA
jbartig@internet2.edu

Sadi Koçak
SURF
Utrecht, Netherlands
sadi.kocak@surf.nl

Niels den Otter
SURF
Utrecht, Netherlands
niels.denotter@surf.nl

Kaj Koole
SURF
Utrecht, Netherlands
kaj.koole@surf.nl

James Deaton
Internet2
Ann Arbor, MI, USA
jed@internet2.edu

k claffy
CAIDA
UC San Diego
La Jolla, CA, USA
kc@caida.org

Abstract

BGP hides information that is crucial for building accurate routing models. In this paper, we combine BGP and active probing to infer relative route preference policies of research and education (R&E) connected ASes. We inferred that systems in ≈88% of ≈12K prefixes that 2,578 ASes announced in the R&E ecosystem were insensitive to AS path length when selecting provider routes – only ≈8-9% appeared to assign the same local preference to available R&E and commodity routes. We validate our method, and discuss broader application of the method to infer relative route preference, a crucial step in being able to accurately model routing policies.

CCS Concepts

• **Networks → Network measurement.**

Keywords

- Comparison of May 2025 (SURF) and June 2025 (Internet2) experiments
- Comparison of localpref inference (egress policy) with AS path prepending observations (ingress policy)
- Case study of route selections when a network assigned equal localpref

ACM Internet Measurement Conference, October 2025

ROOTBEER project

Routing Operations Observational Technology: Building to Enable Education and Research

- We built a prototype to investigate R&E localpref routing policies that was difficult to operationalize
- We recognized that there are security-relevant properties of R&E routing that we could now investigate
- We are now building an operational system to support R&E routing policy inferences, hosted by Internet2
- NSF CICI OAC-2530871

ROOTBEER theme: identifying problematic routing

R&E networks generally prefer routes from other R&E networks, so **R&E route announcements are high-impact.**

1. A member **leaking commodity routes into R&E** can attract traffic they might not be prepared for.

2. A member **leaking R&E routes into commodity** can result in traffic across R&E infrastructure that should use commodity.

3. A member announcing an **anycast prefix into R&E** can attract traffic best directed to other anycast instances.

These problems are not theoretical

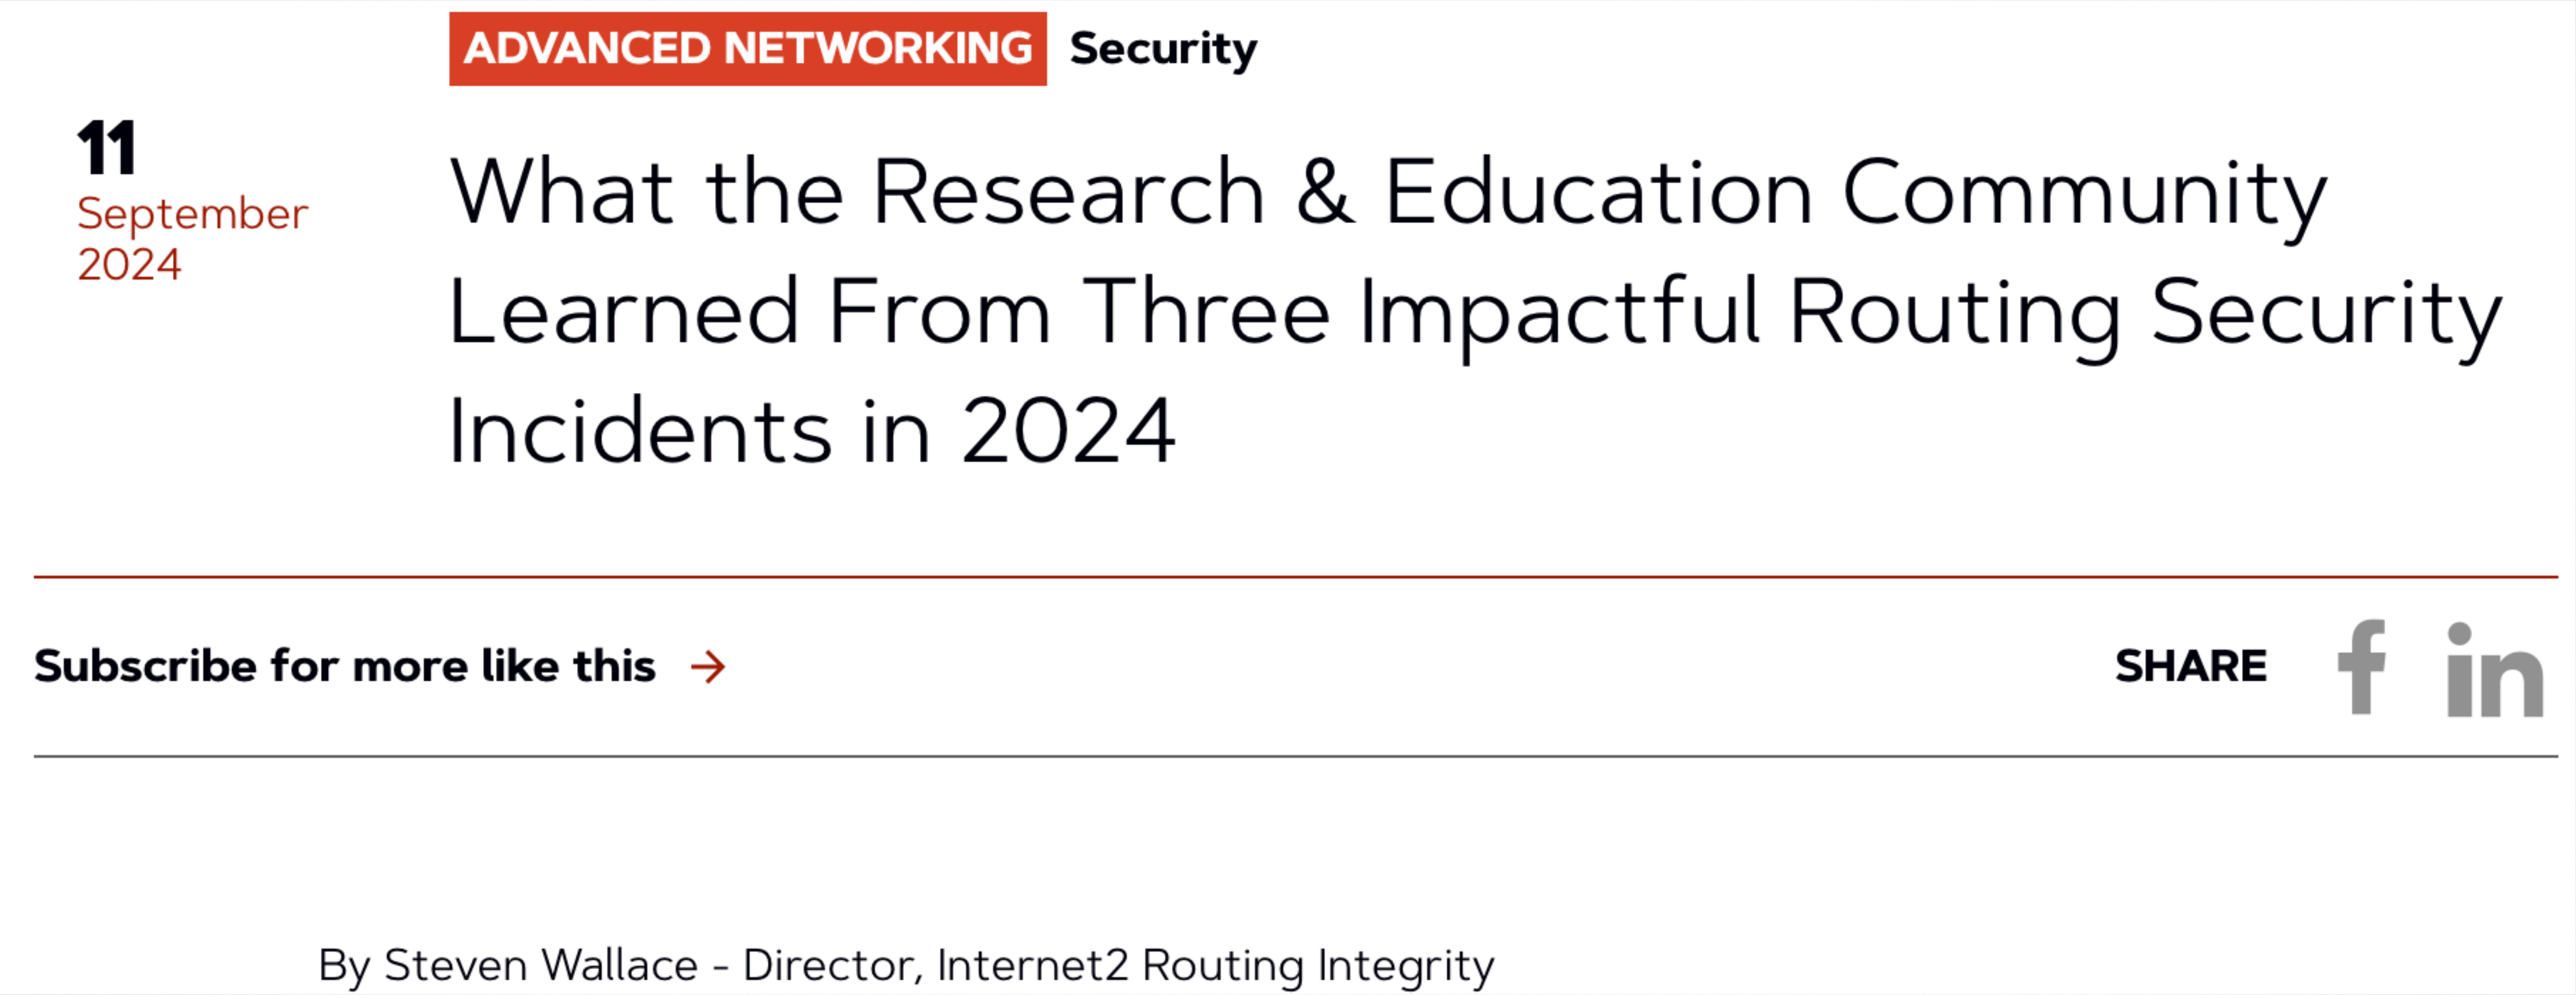

ADVANCED NETWORKING Security

11
September
2024

What the Research & Education Community Learned From Three Impactful Routing Security Incidents in 2024

Subscribe for more like this →

SHARE

By Steven Wallace - Director, Internet2 Routing Integrity

These problems are not theoretical

- **Incident #1: Commercial Routes Leaked to R&E**
  - In March 2024, a South American R&E network mistakenly announced commercial internet routes to its R&E peers. Other R&E networks used these routes, causing traffic to be rerouted through South America.
- **Incident #2: R&E Routes Leaked to Commercial Providers**
  - In August 2024, a R&E member in the Middle East accidentally leaked routes to its commercial Internet provider. For networks announcing more specific routes to R&E, this misconfiguration caused returning traffic from major cloud providers to route through the Middle East.
  - Some members have withdrawn more specific R&E routes to minimize risks.

A Microsoft IPv6 leak by a network in Asia, then reannounced to CERNET (CERNET traffic crossing the Pacific twice)

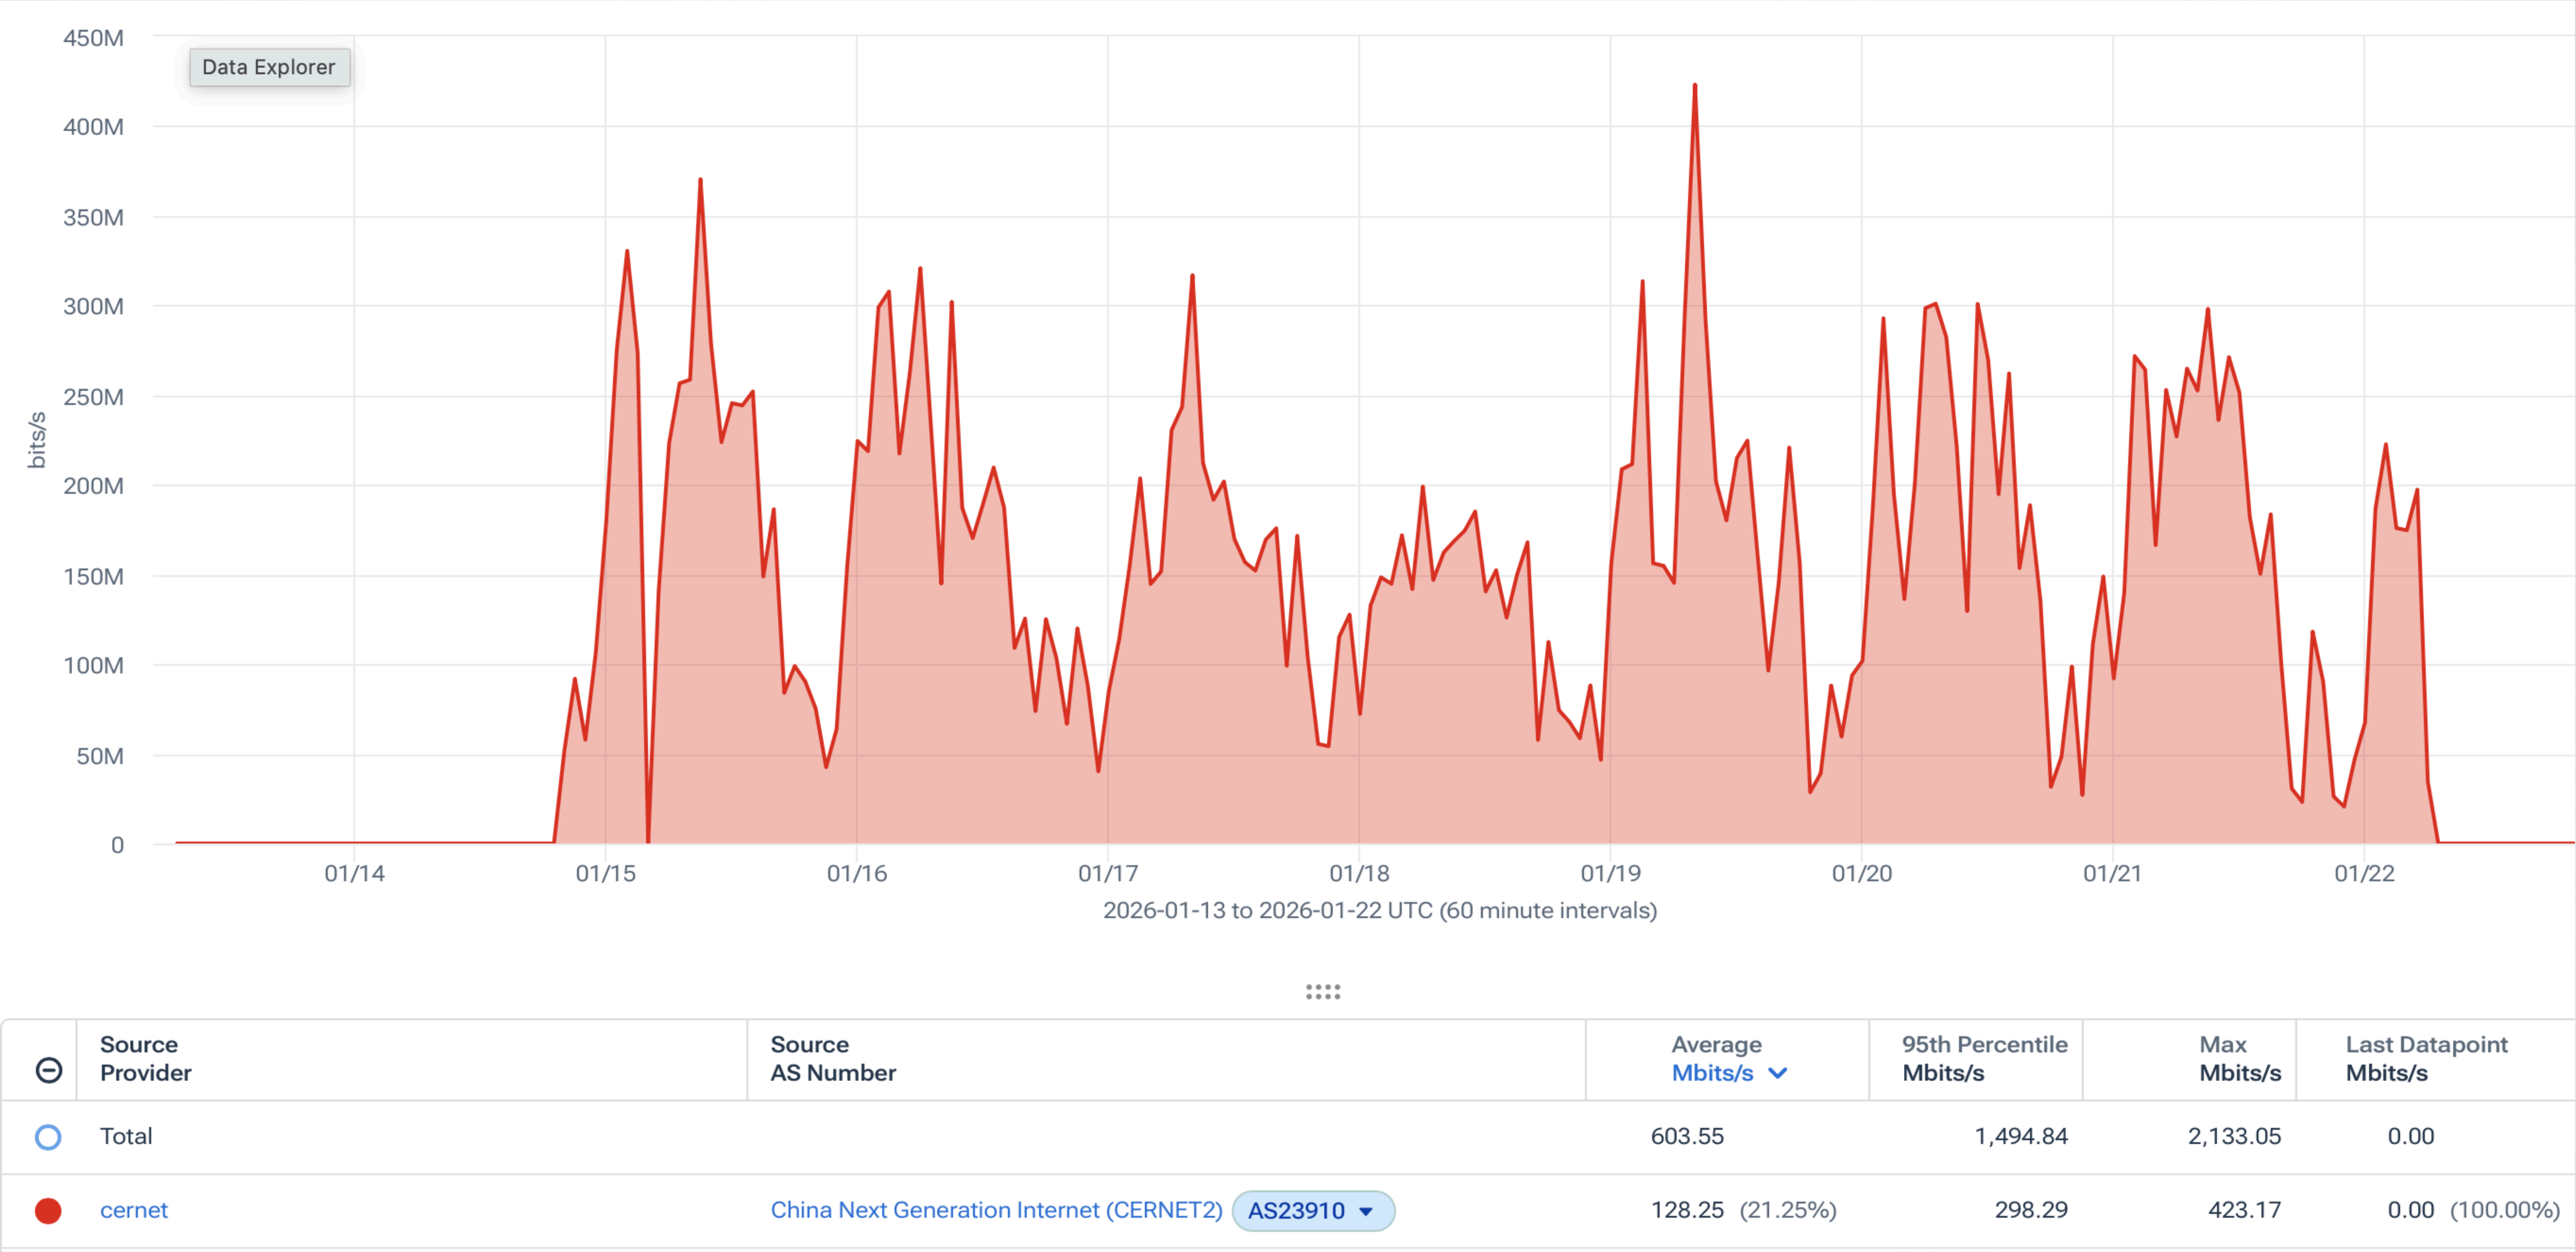

⊖	Source Provider	Source AS Number	Average Mbits/s ⌄	95th Percentile Mbits/s	Max Mbits/s	Last Datapoint Mbits/s
○	Total		603.55	1,494.84	2,133.05	0.00
●	cernet	China Next Generation Internet (CERNET2) AS23910 ▾	128.25 (21.25%)	298.29	423.17	0.00 (100.00%)

Towards Preventing Leaked Routes

- Because R&E route announcements have high impact in R&E networking, Internet2 filters route announcements:
  - U.S. R&E peers: members authorize Internet2 to announce their prefixes
  - International R&E peers: light filtering (bogons, Tier-1s in path)
- Plan is to move towards an ***actively maintained*** public filter on all peers
  - Allows for current mutual transit arrangements while preventing leaks
  - List can be shared and used by other R&E transit networks
- ROOTBEER: building systems support, NSF-funded grant OAC-2530871

Investigating R&E Routing Policies at Scale

Requirements

1. Ability to originate routes whose properties we'll measure
2. Measure global effects of route announcements without requiring measured ASes to do anything to assist the inference.

Instead of probing just R&E-advertised prefixes, we can use the same idea and probe all routed prefixes

ROOTBEER theme: identifying problematic routing

R&E networks generally prefer routes from other R&E networks, so **R&E route announcements are high-impact.**

1. A member **leaking commodity routes into R&E** can attract traffic they might not be prepared for.

Leaking **into** R&E

Internet2 configures prefix filters, accepting routes that match prefixes previously authorized by operators

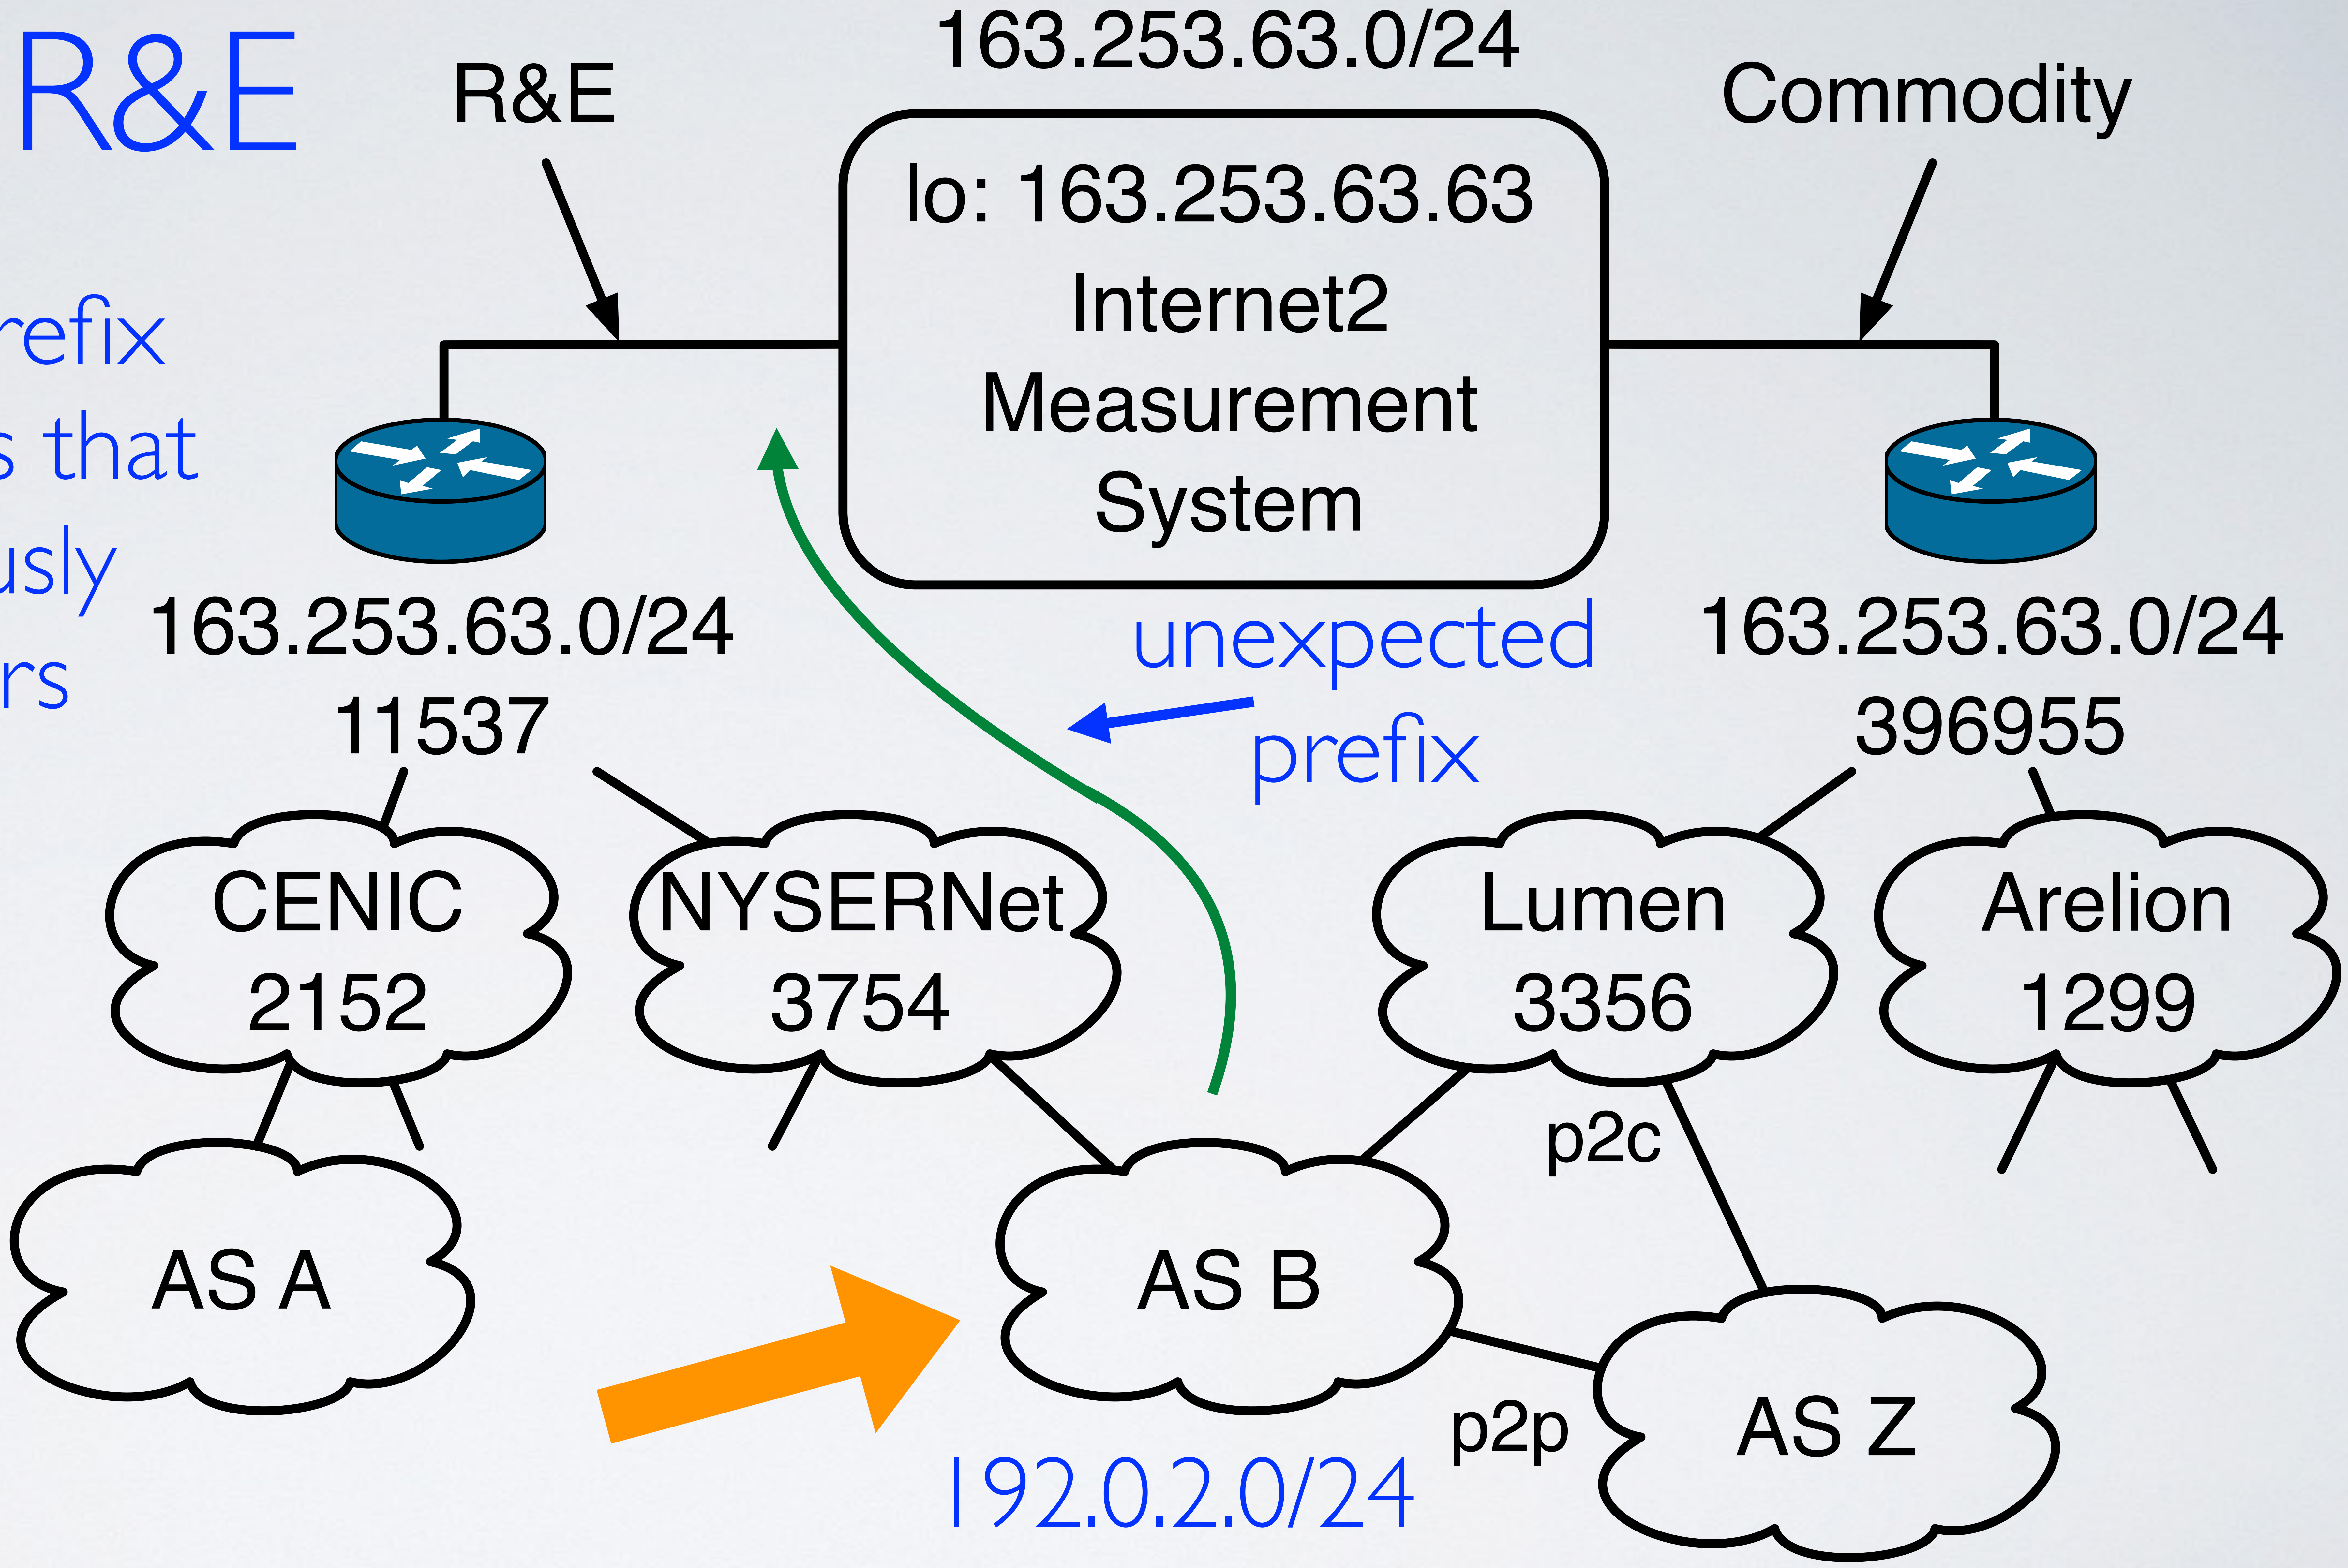

How do systems in the unexpected prefix reach Internet2?

Leaking **into** R&E

Internet2 configures prefix filters, accepting routes that match prefixes previously authorized by operators

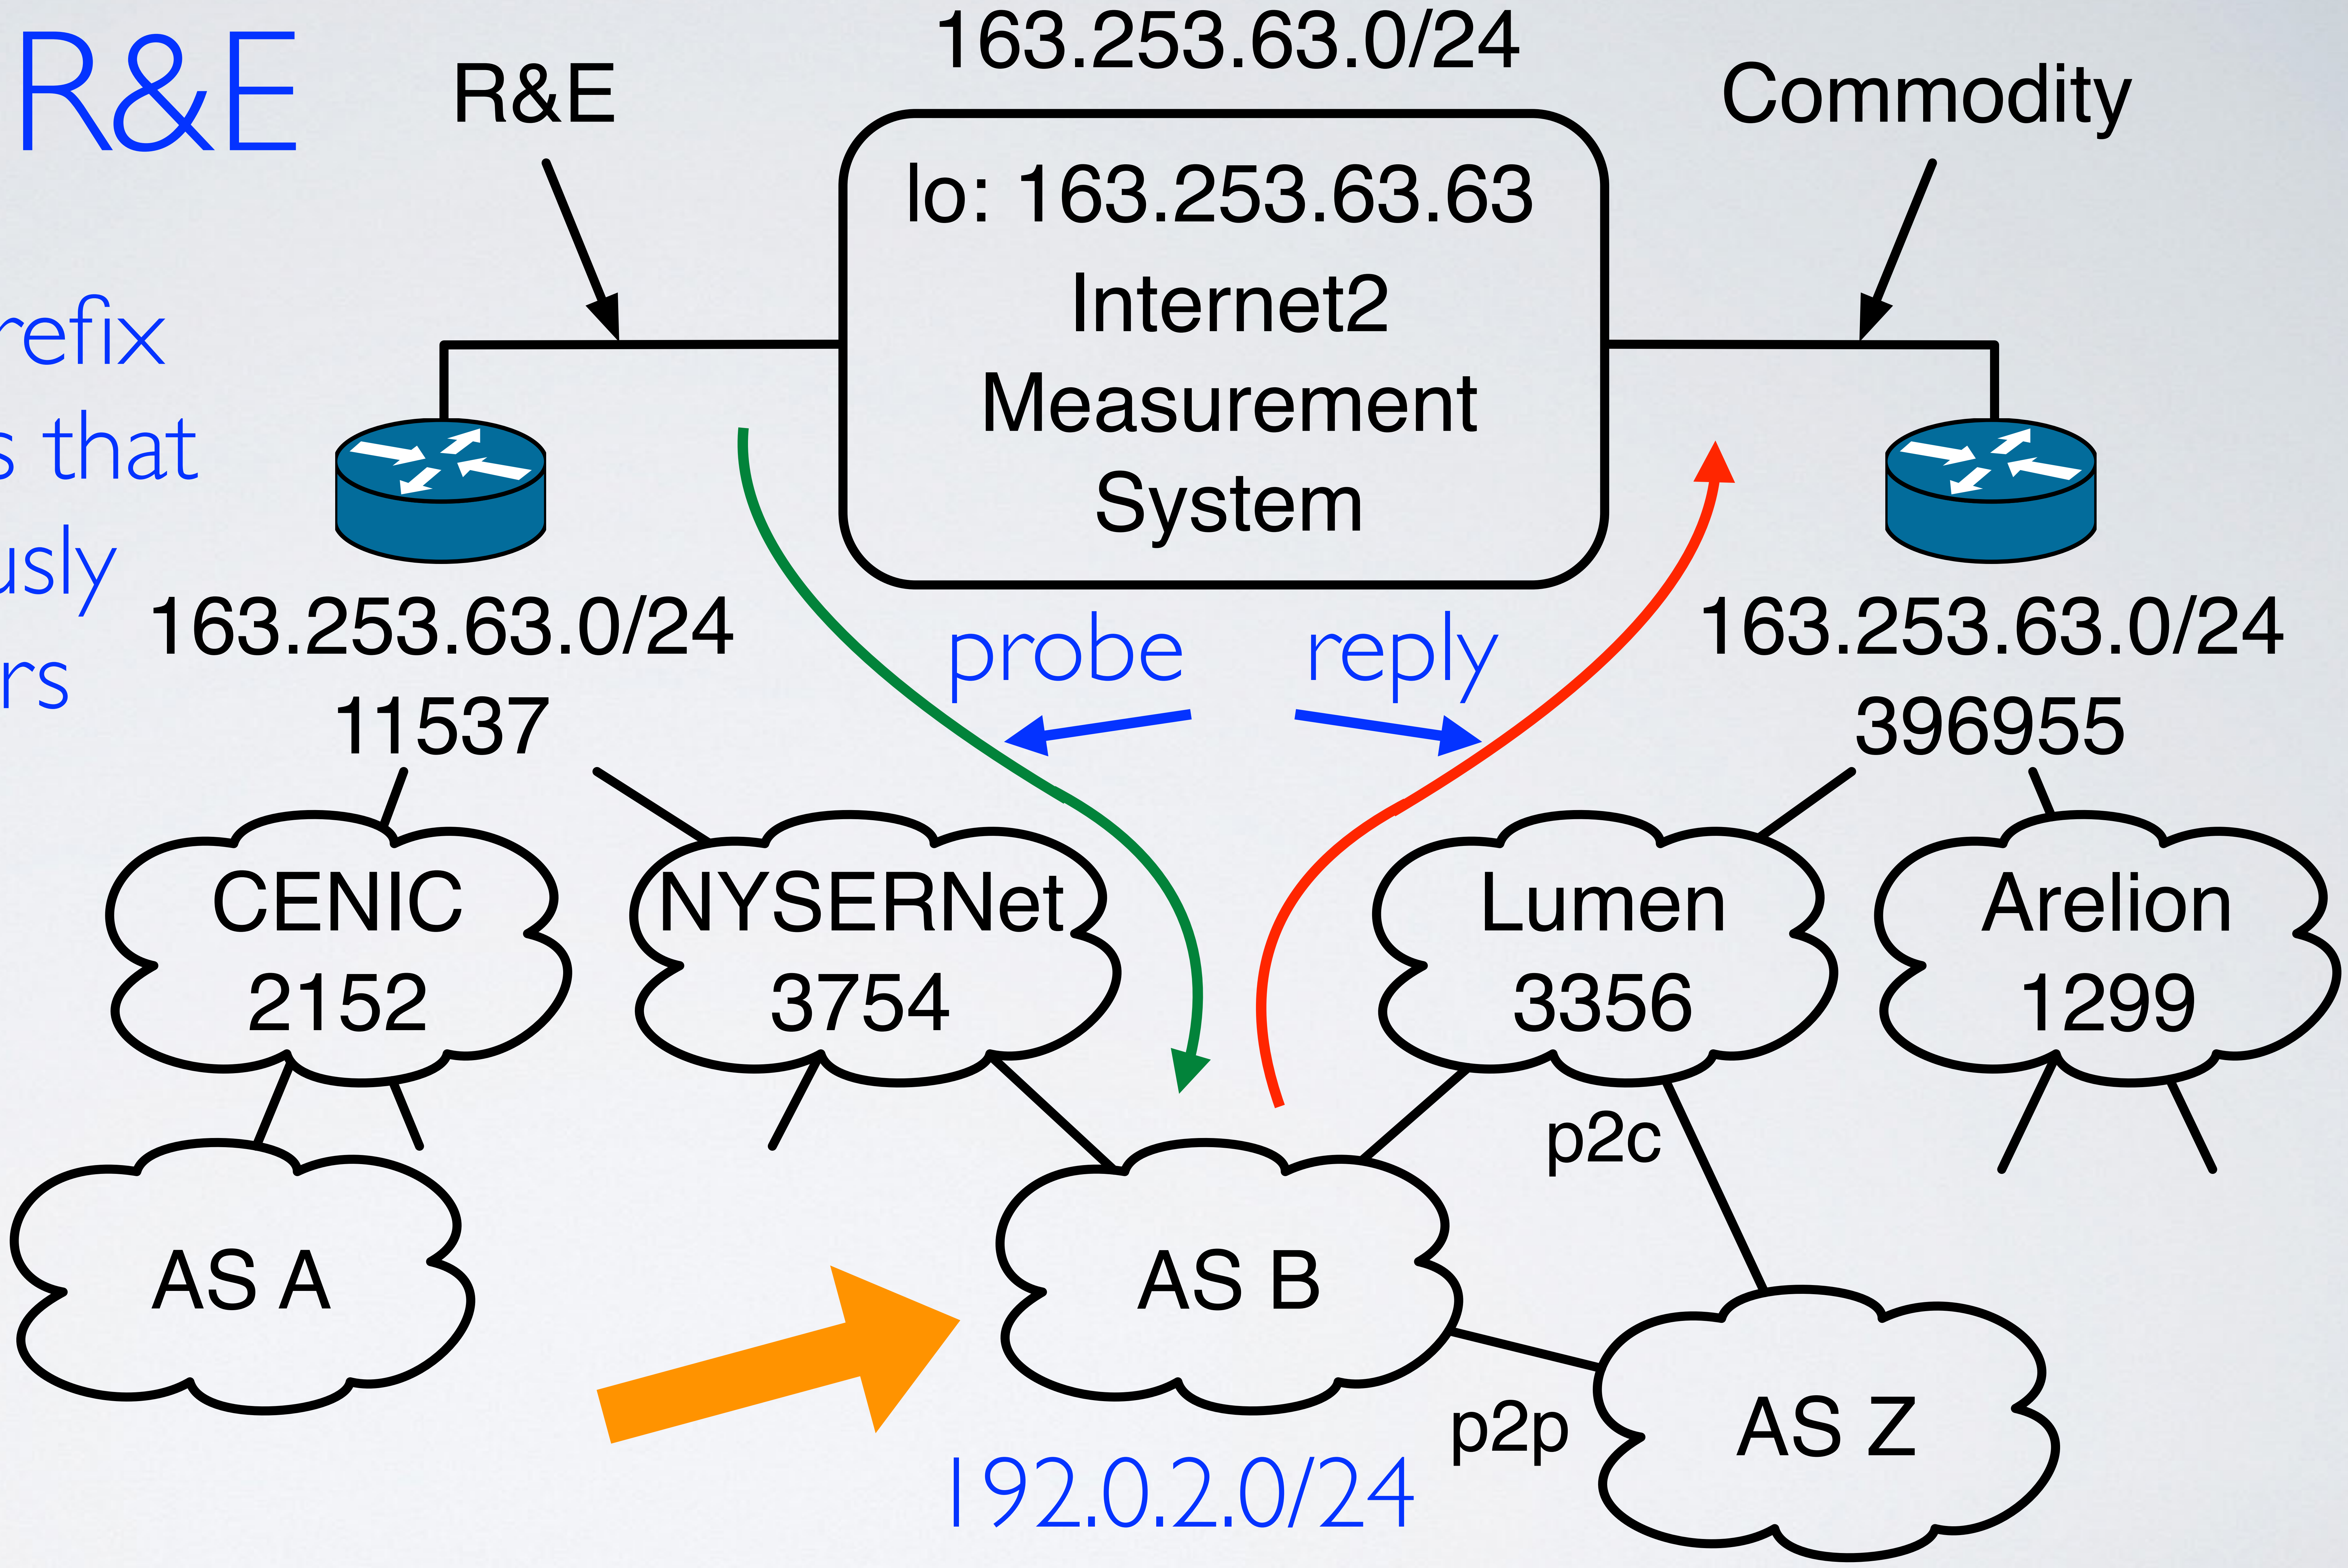

Responses via commodity may indicate B's prefix was leaked

Leaking into R&E: initial results

- As of 24 Nov 2025

- 145 IPv4 prefixes not covered by other previously authorized R&E routes were filtered by Internet2

- 54 (37%) of these prefixes responded via commodity - a rate 5x higher than authorized R&E routes

- Originated by 31 ASes, supplied by 8 peers.

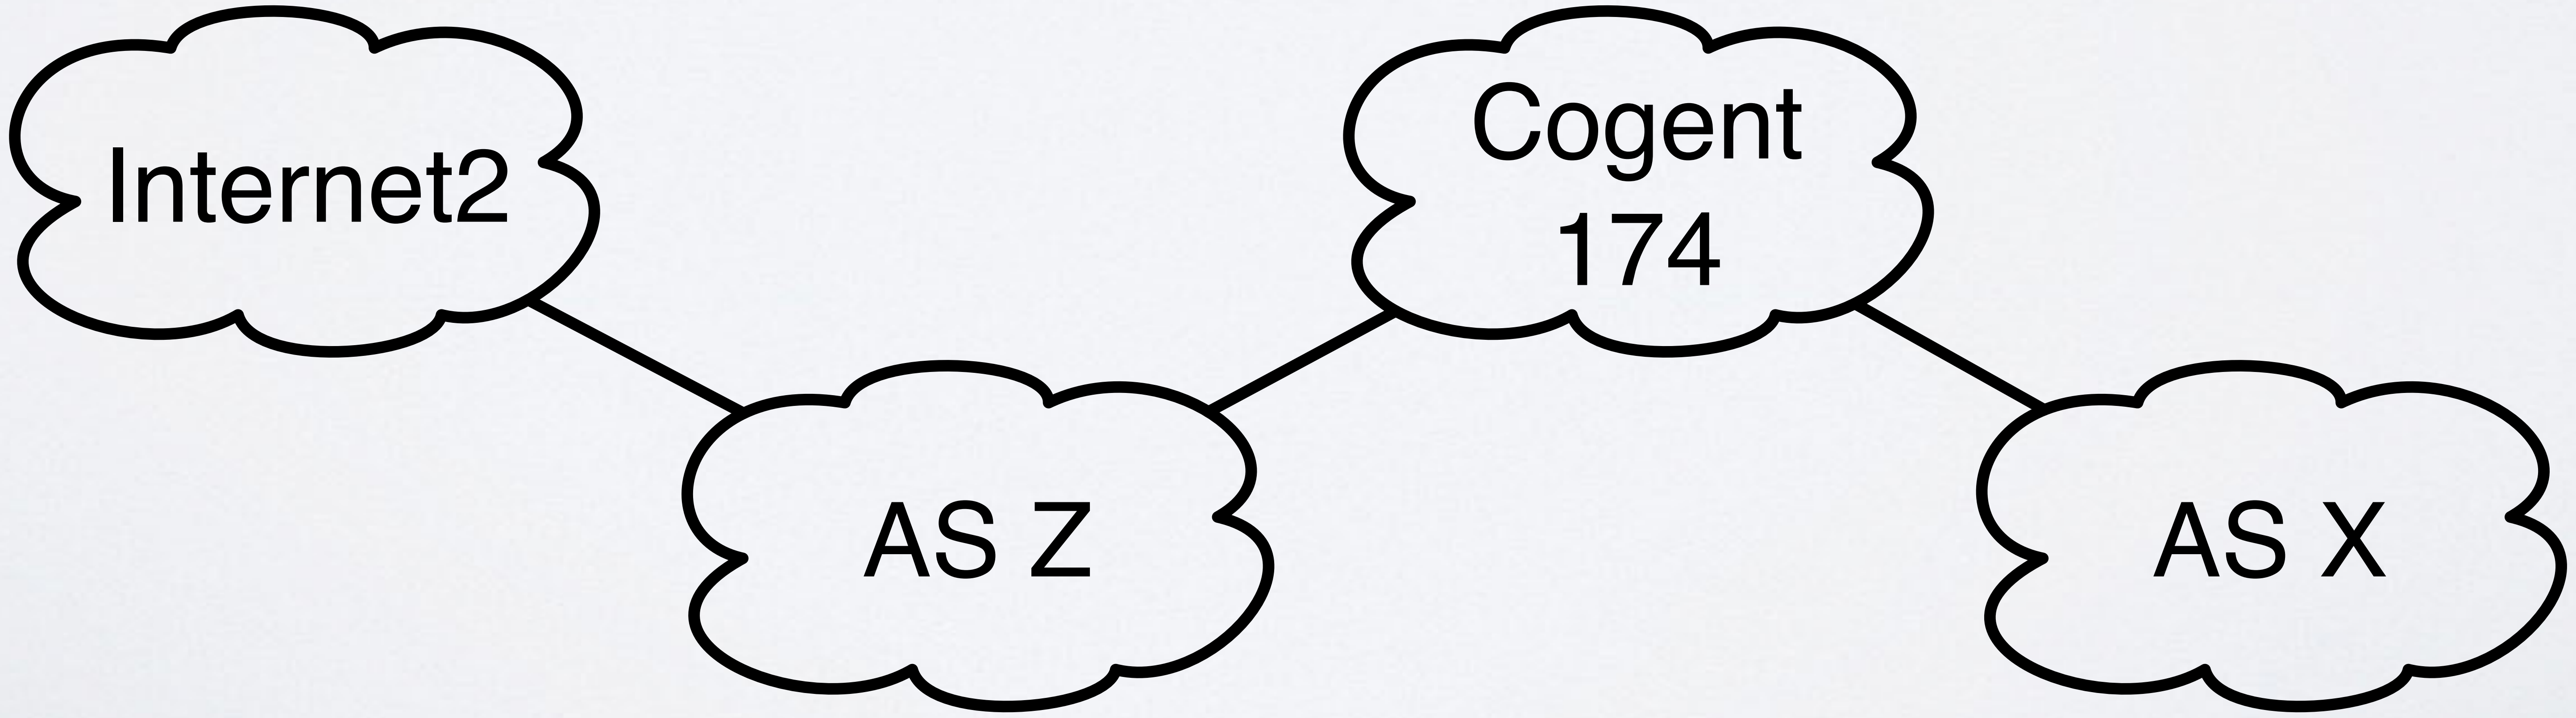

ROOTBEER theme: identifying problematic routing

R&E networks generally prefer routes from other R&E networks, so **R&E route announcements are high-impact.**

1. A member **leaking commodity routes into R&E** can attract traffic they might not be prepared for.

2. A member **leaking R&E routes into commodity** can result in traffic across R&E infrastructure that should use commodity.

Leaking **from** R&E

If R&E network B shares routes with commodity network Z

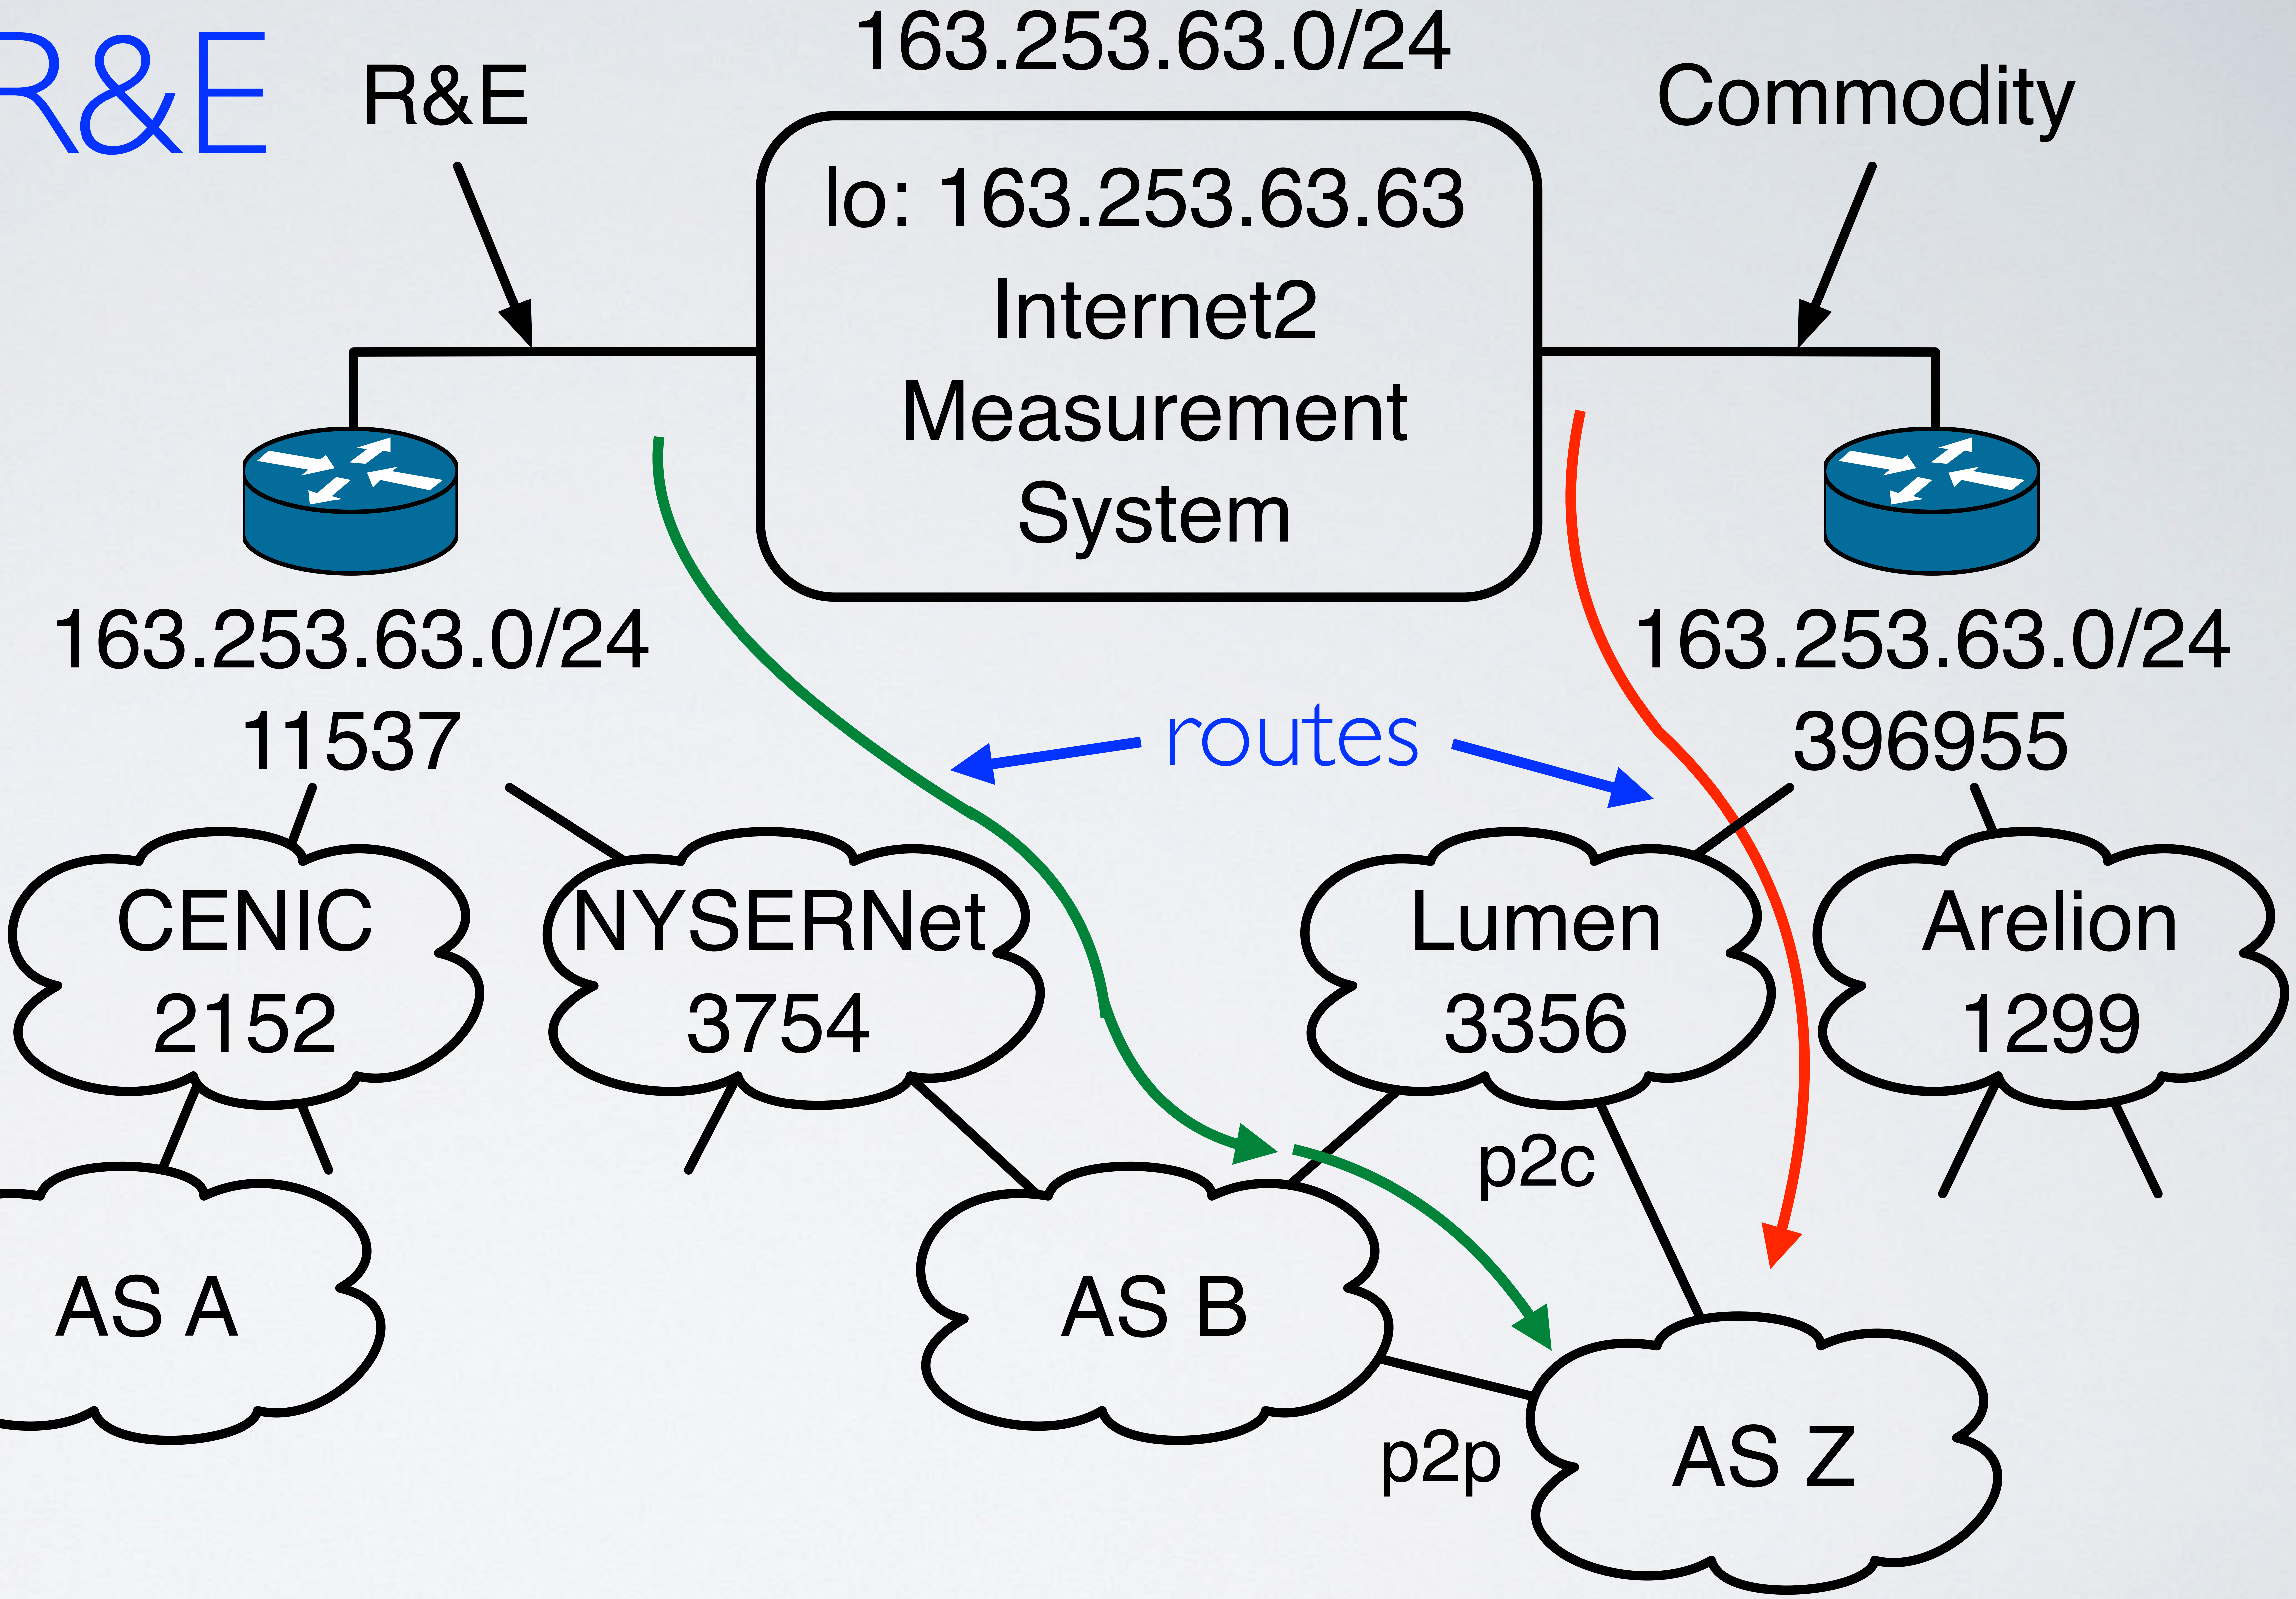

From Z: Leaked R&E route: `B 3754 11537 163.253.63.0/24`

Commodity route: `3356 396955 163.253.63.0/24`

Leaking **from** R&E

A ping probe to a responsive destination in Z ….

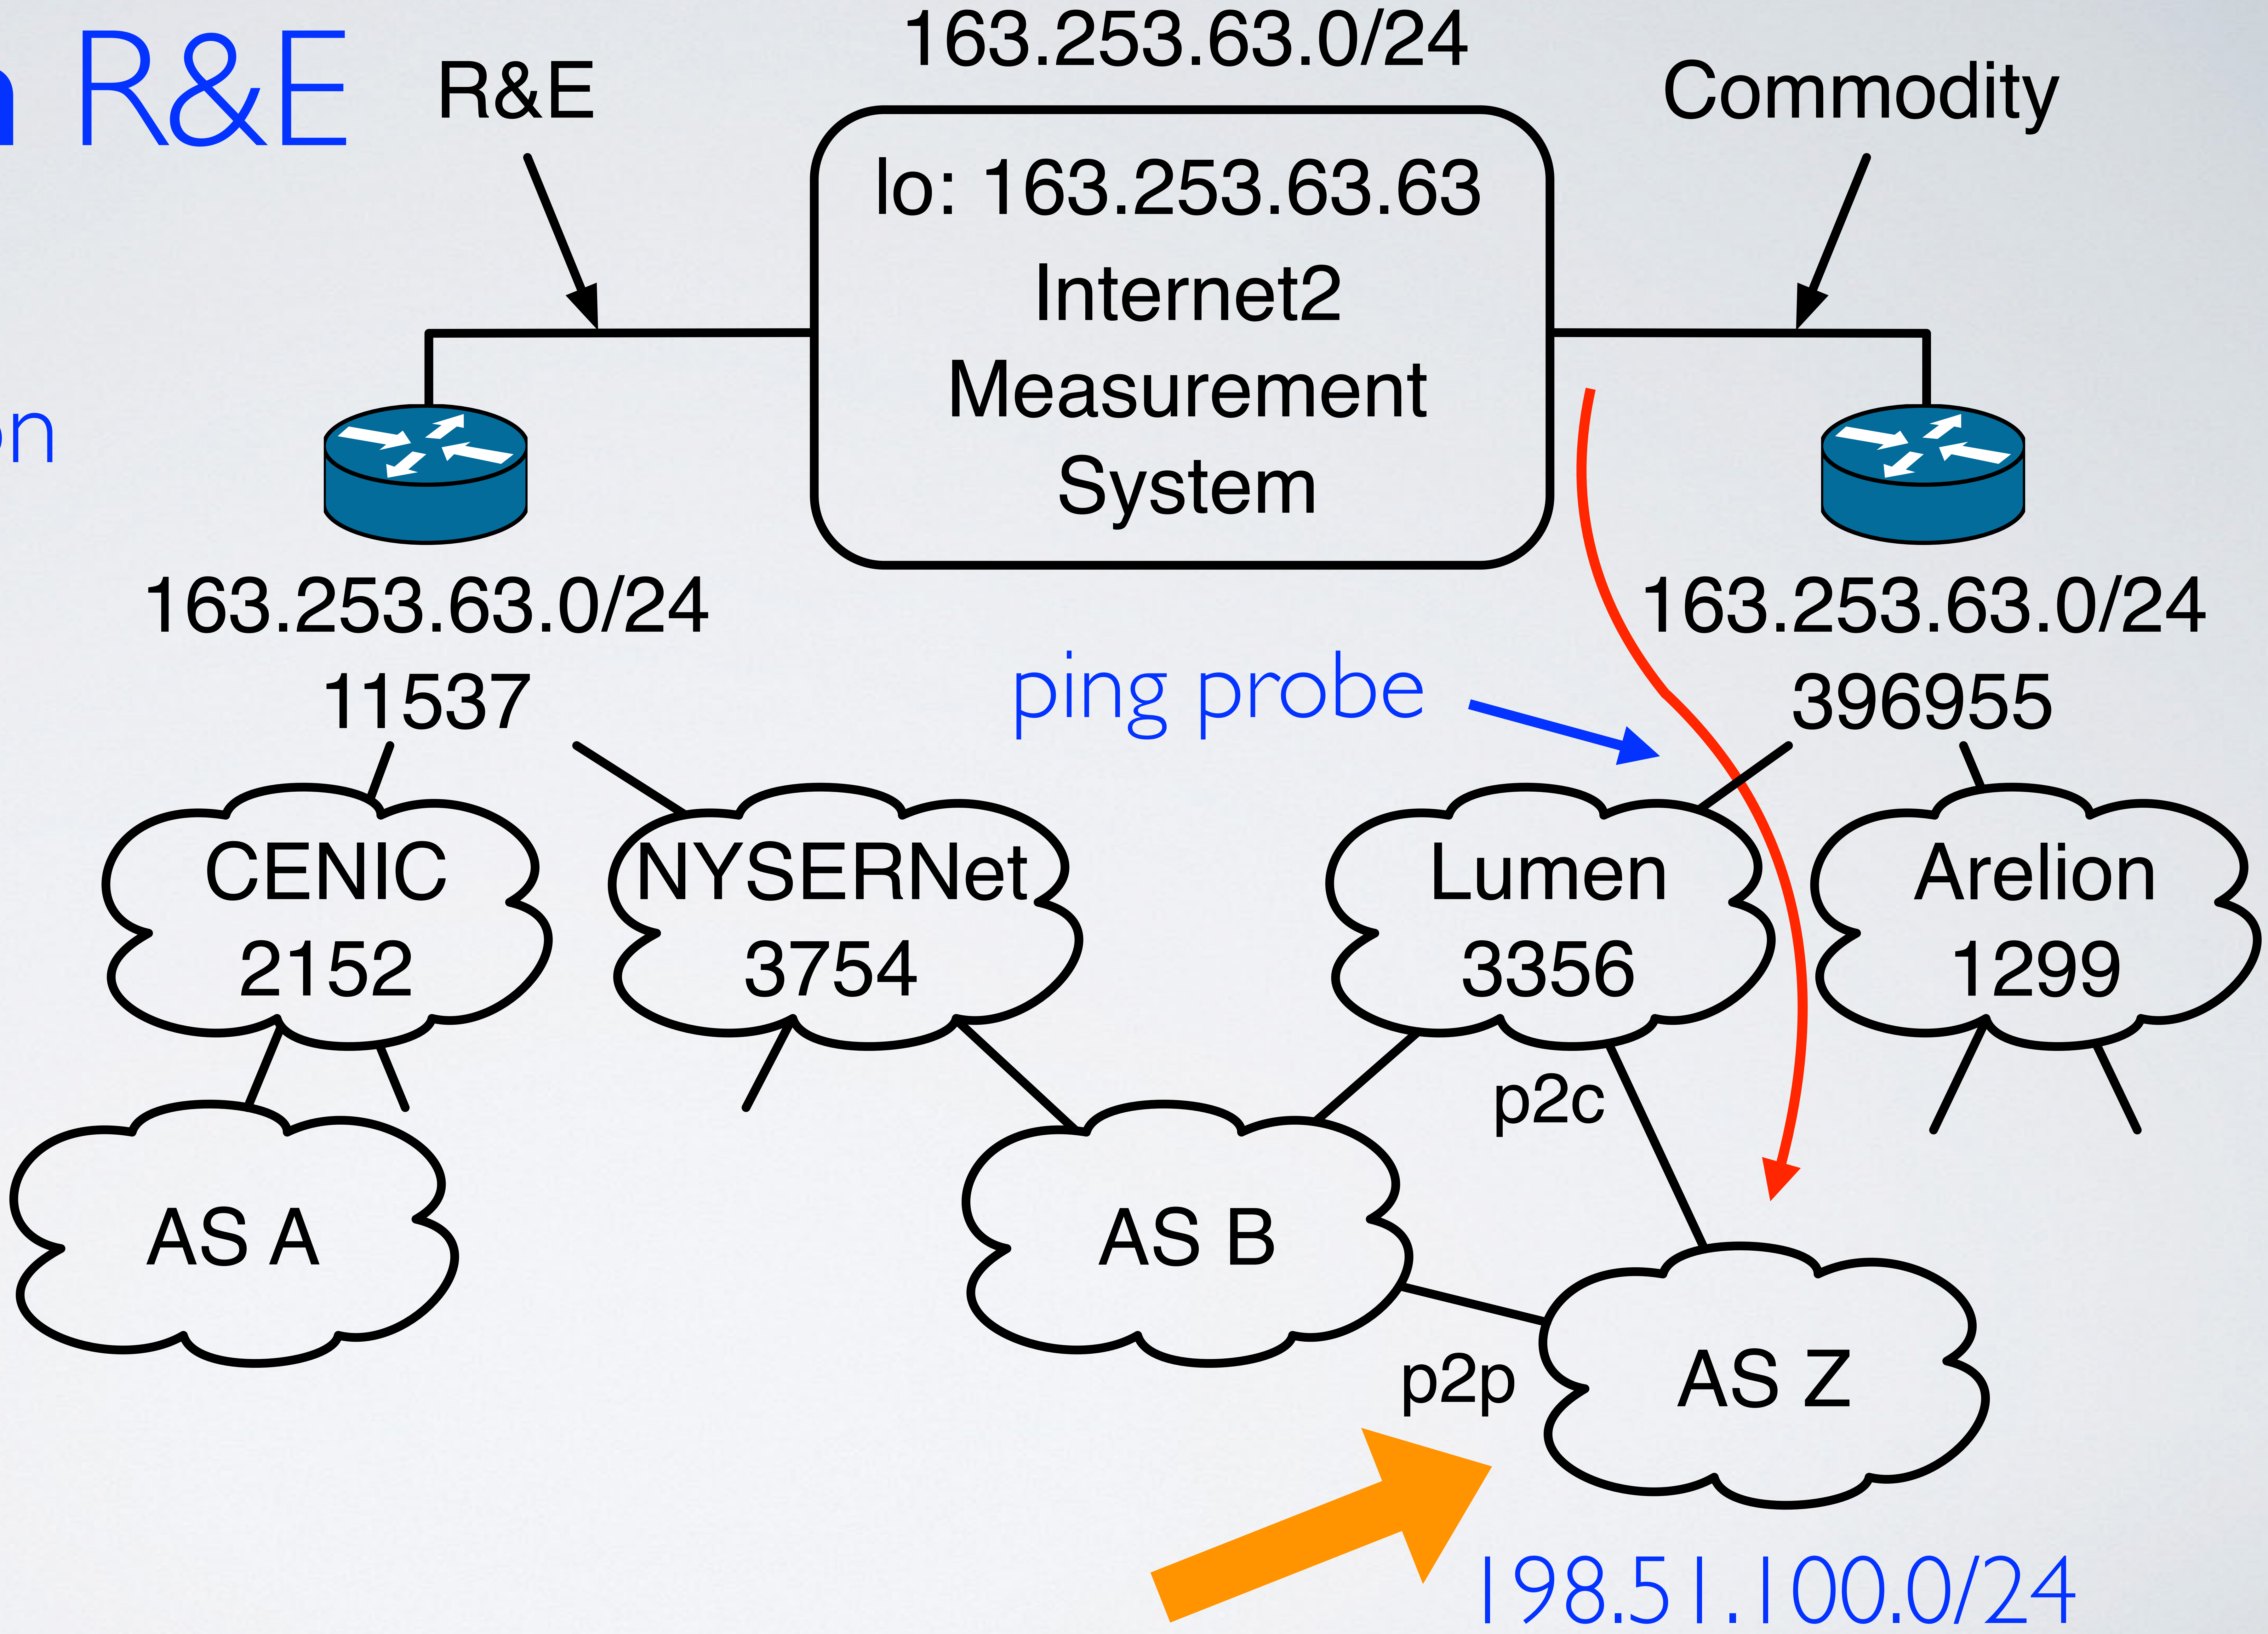

Leaking **from** R&E

A ping probe to a responsive destination in Z ….

… follows an R&E path back to our measurement system…

…it may be leak

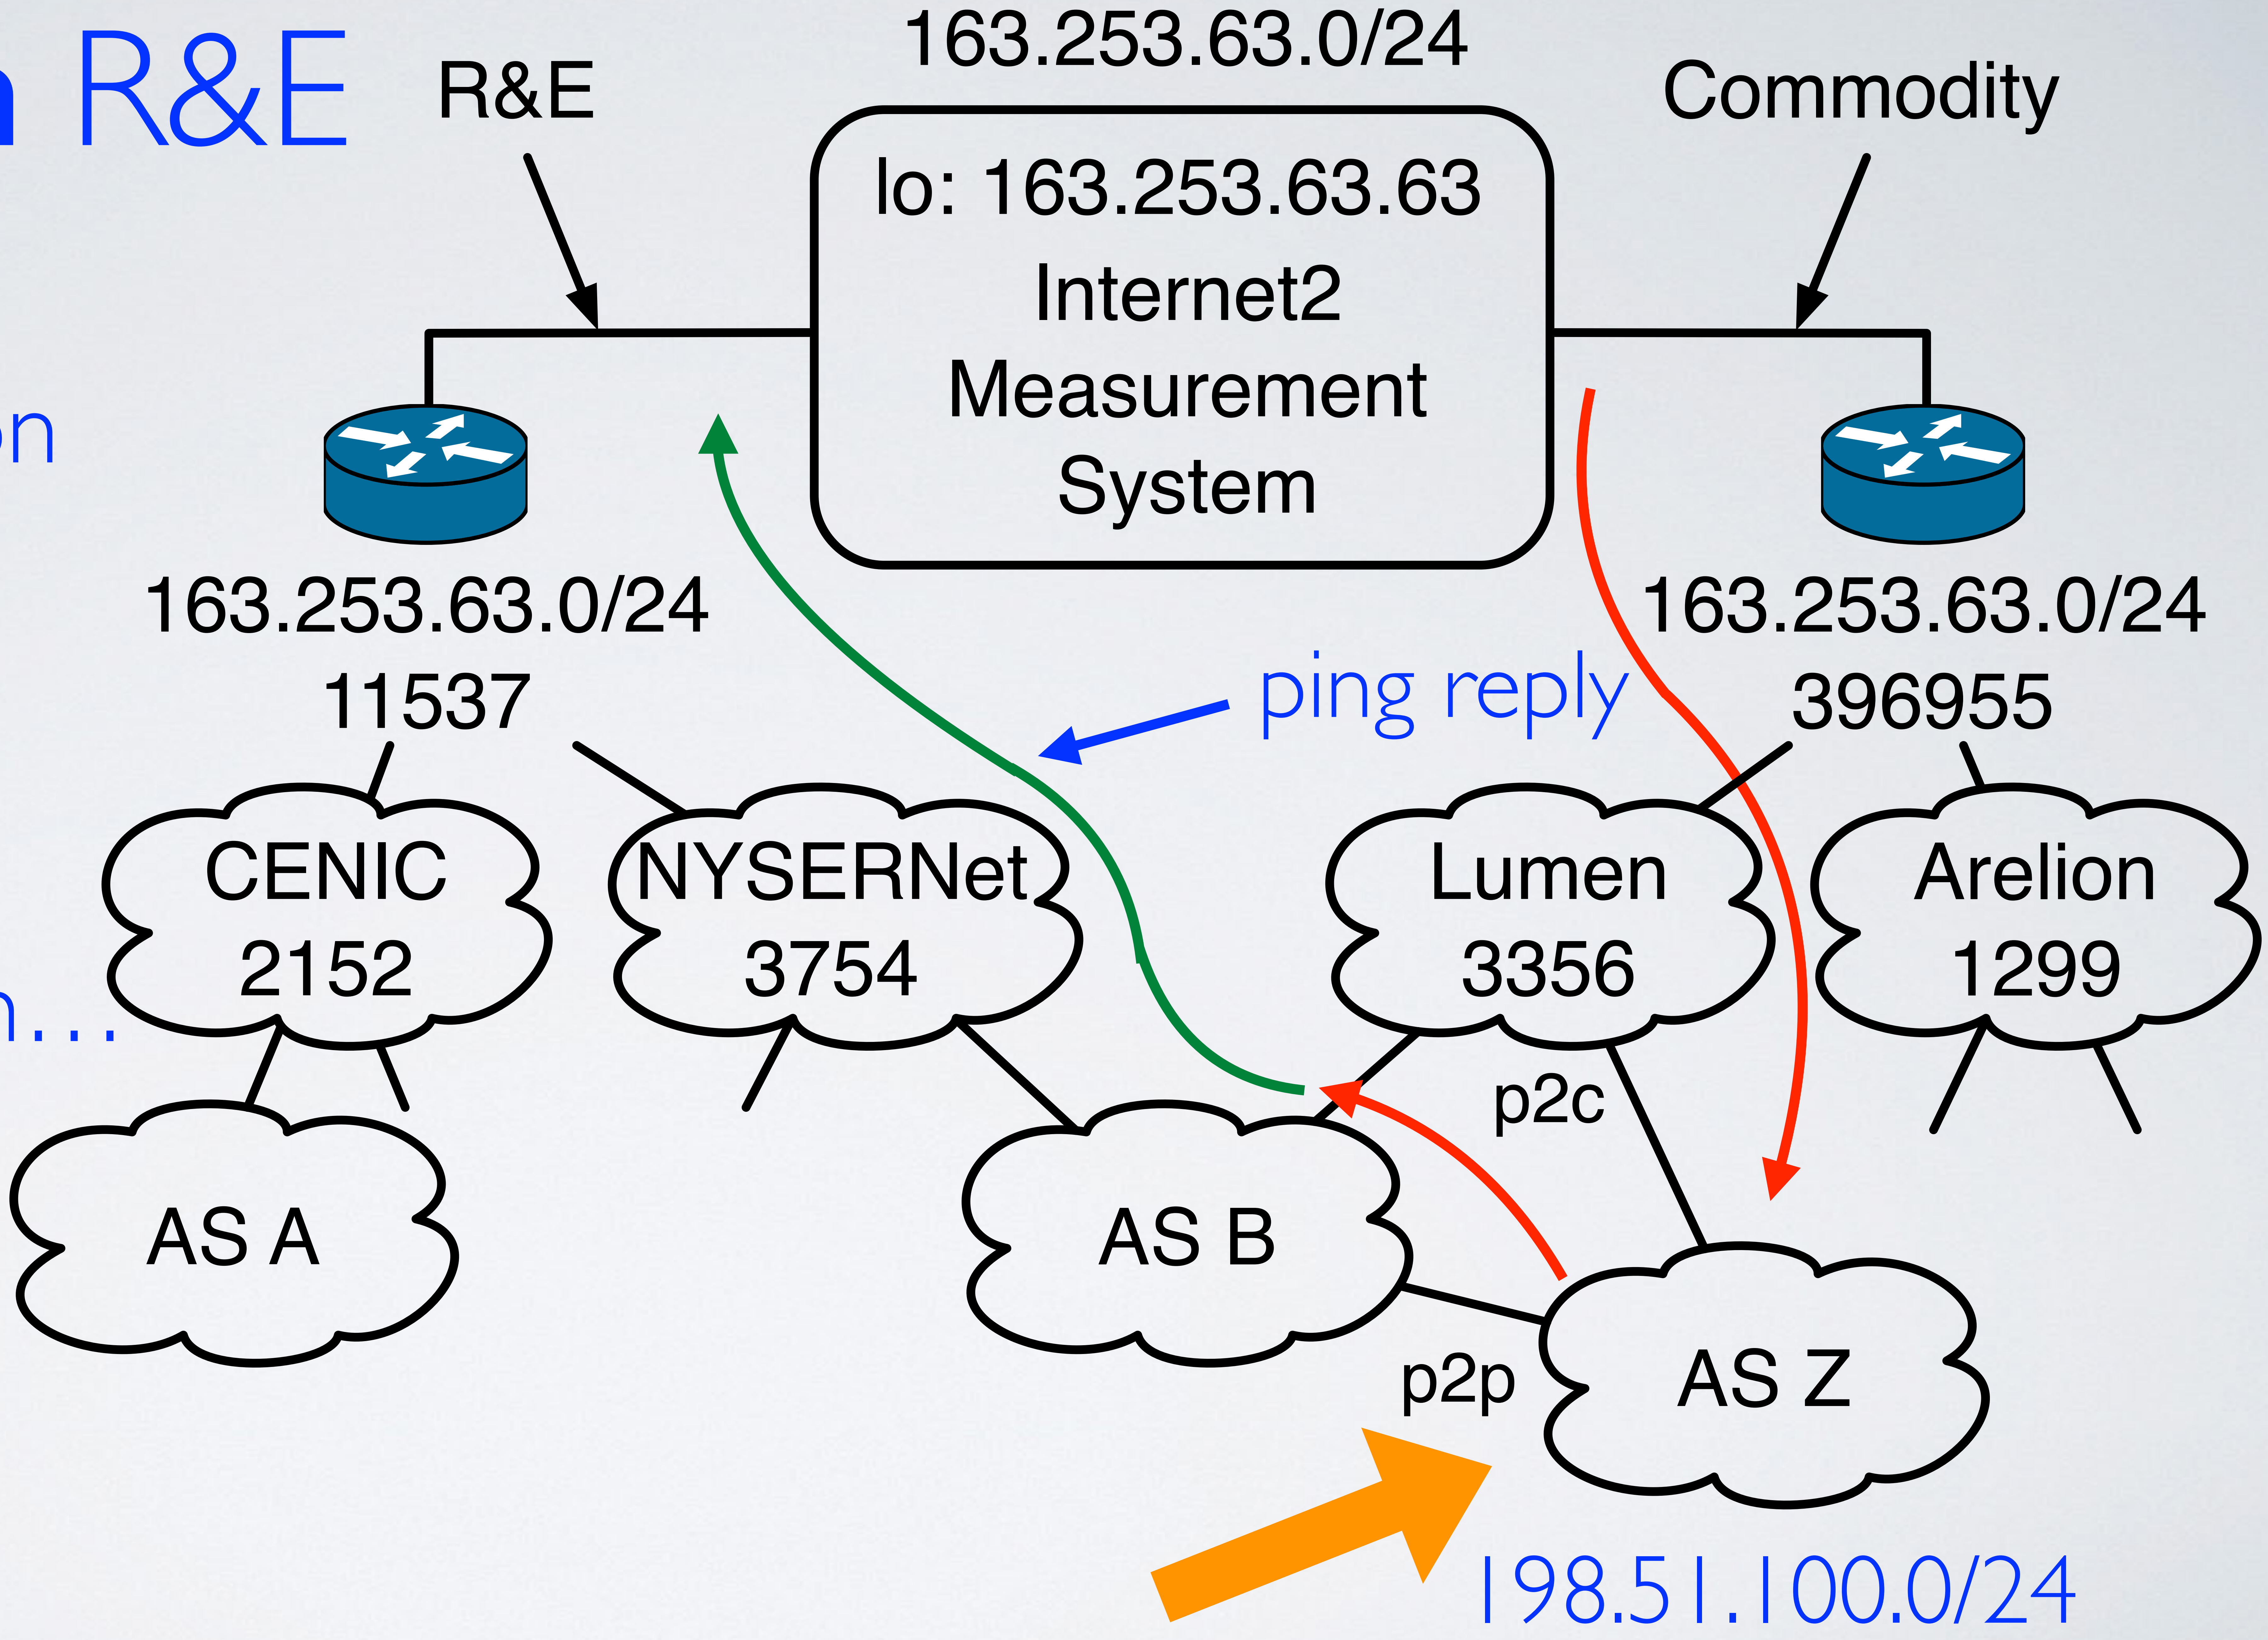

ROOTBEER theme: identifying problematic routing

R&E networks generally prefer routes from other R&E networks, so **R&E route announcements are high-impact.**

1. A member **leaking commodity routes into R&E** can attract traffic they might not be prepared for.

2. A member **leaking R&E routes into commodity** can result in traffic across R&E infrastructure that should use commodity.

3. A member announcing an **anycast prefix into R&E** can attract traffic best directed to other anycast instances.

Measuring Anycast routes in R&E

- Daily anycast census publicly provided by University of Twente: https://github.com/ut-dacs/anycast-census
- We probed responsive addresses in 13K anycast prefixes from ~260 CAIDA Ark VPs. These VPs are in a mix of R&E and commodity networks
- We focused on ~70 anycast prefixes that are also in an R&E routing table reported by Internet2 to RouteViews

Anycast route: APNIC AU DNS server

Internet2 learns route to APNIC anycast services from AARnet, most R&E paths go to Sydney

Commodity 50%: ~50ms
R&E 50%: ~200ms

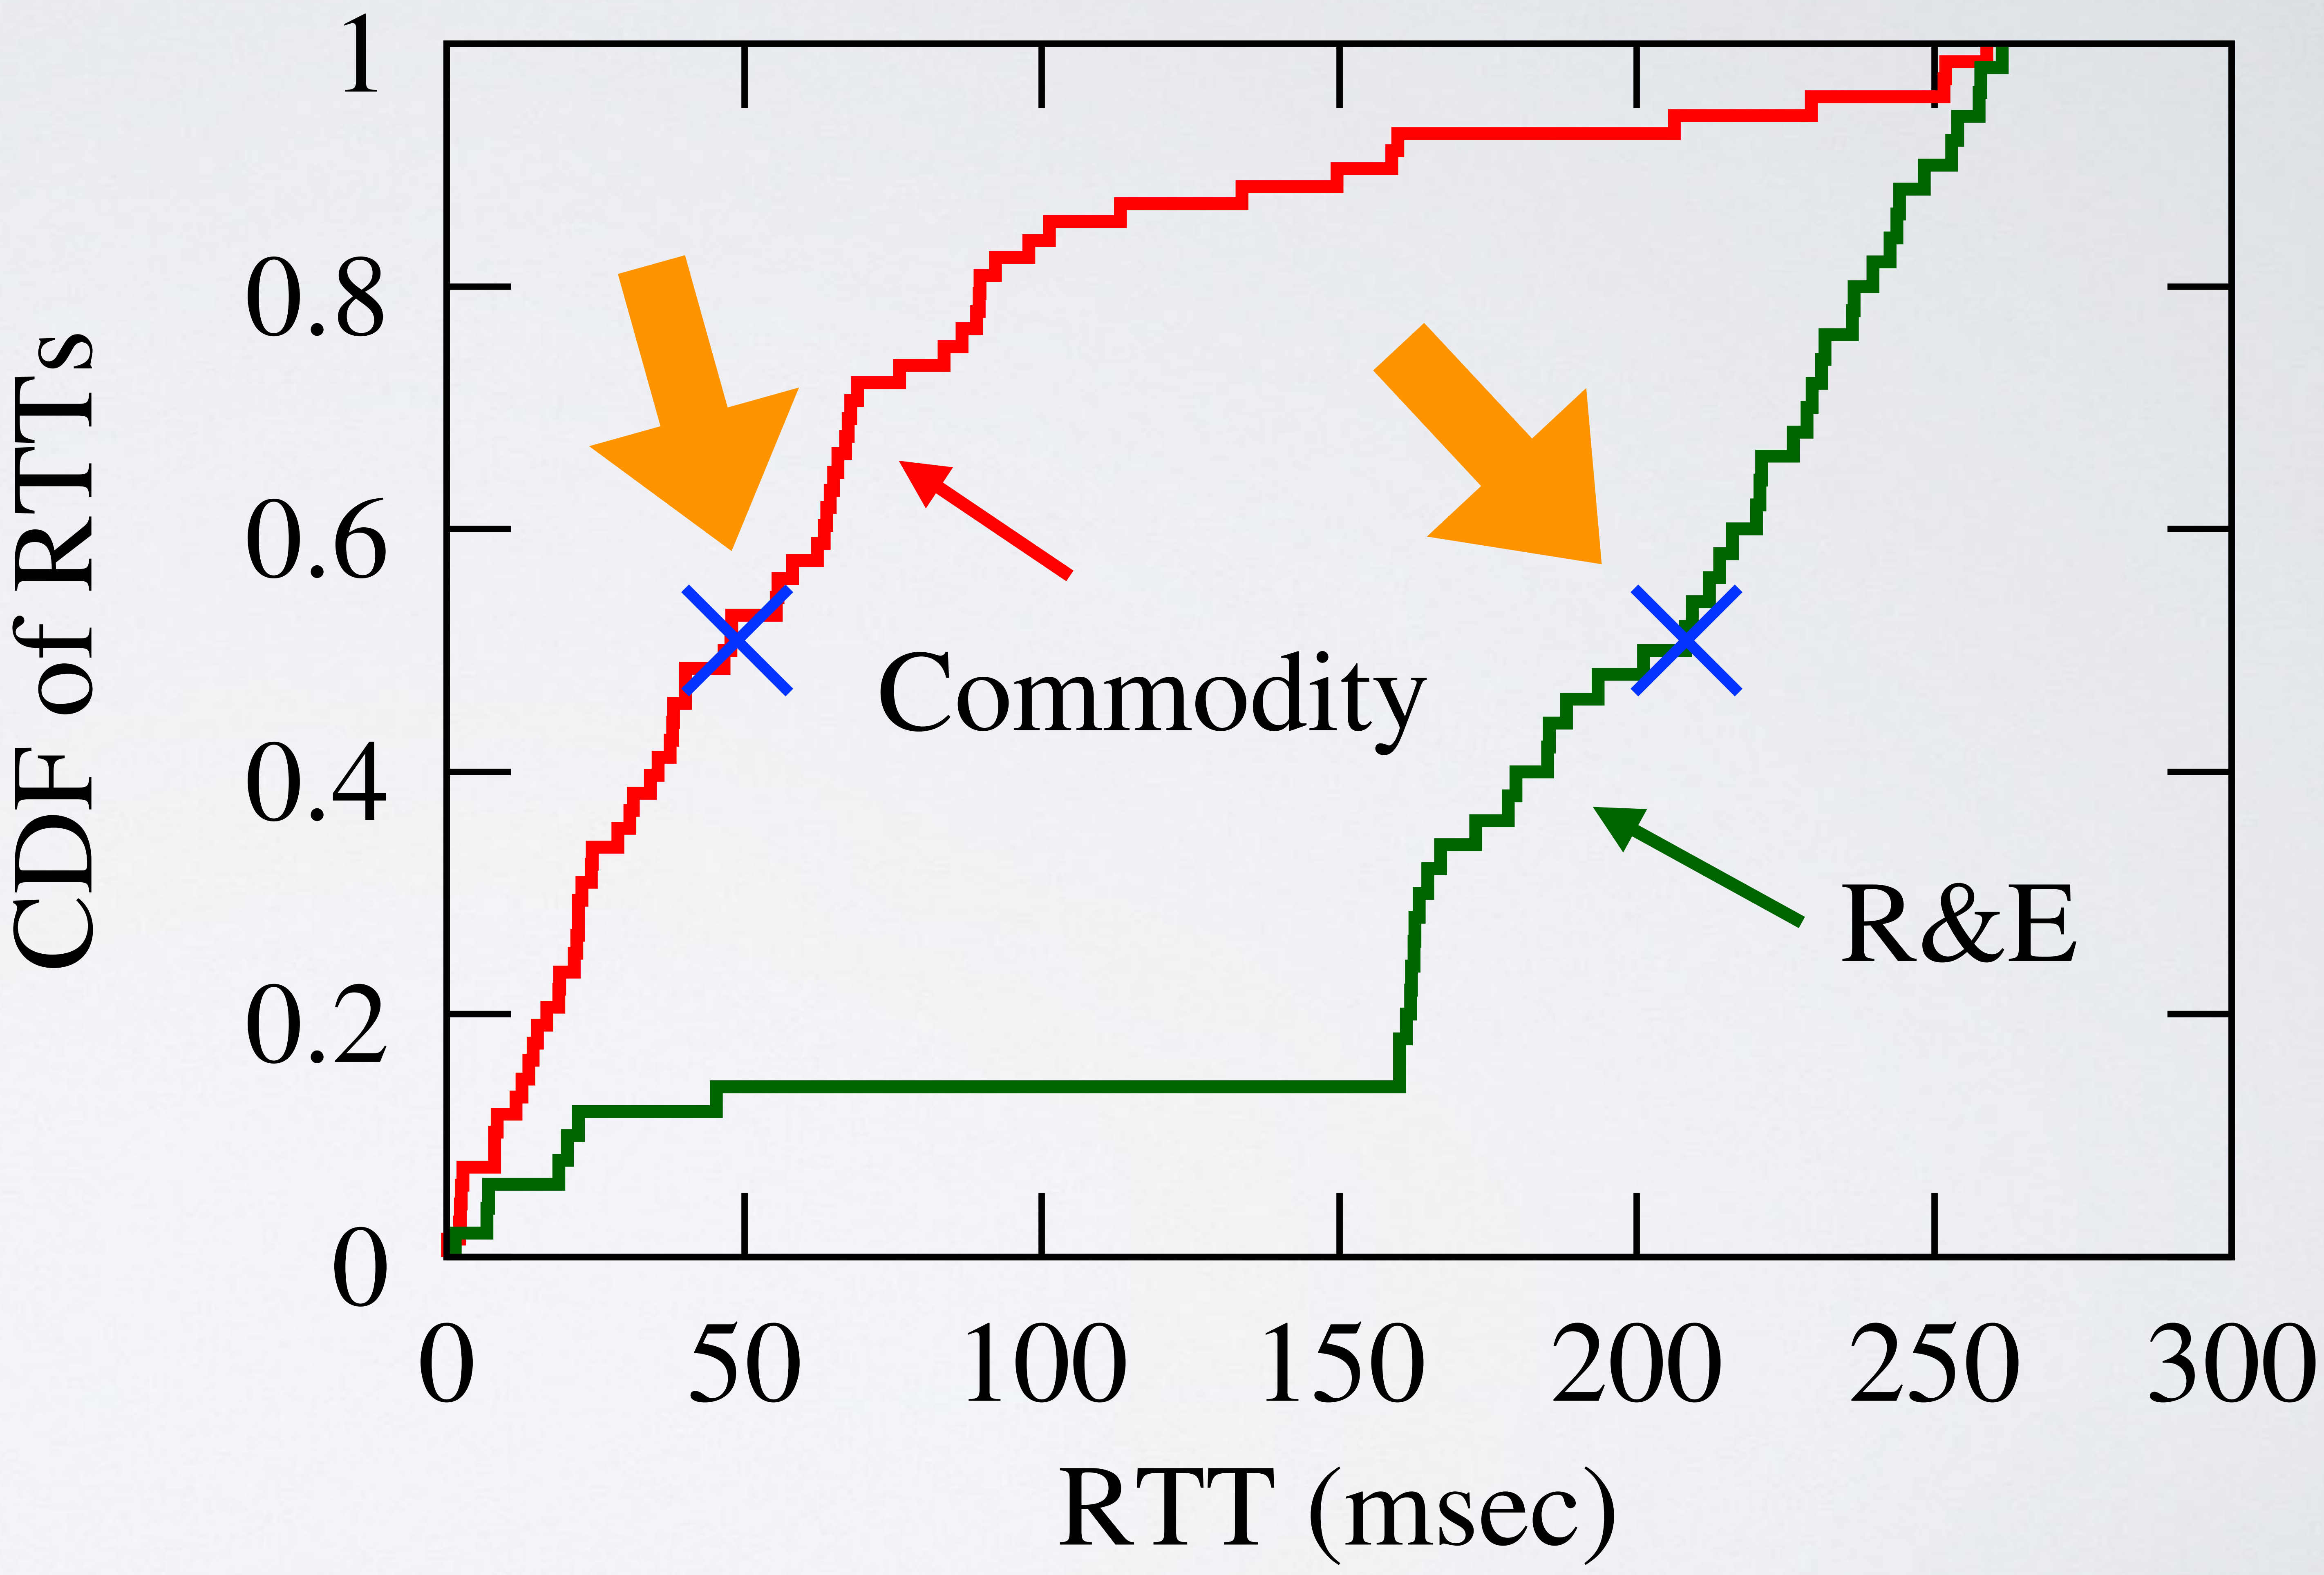

Instance available on U.S. East Coast

Summary

- R&E routing is unique because R&E networks assign higher localpref to R&E routes, and R&E networks provide mutual transit
- ROOTBEER is building a set of operational tools to help with problems.
- Delivered RPKI ROA planning tool already, and we're going to be delivering tools to assist in inferring and improving R&E routing policies.