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

Internet routing, Border Gateway Protocol (BGP)

ACM Reference Format:

Matthew Luckie, Steven Wallace, Karl Newell, Jeff Bartig, Sadi Koçak, Niels den Otter, Kaj Koole, James Deaton, and k claffy. 2025. R&E Routing Policy: Inference and Implication. In *Proceedings of the 2025 ACM Internet Measurement Conference (IMC '25), October 28–31, 2025, Madison, WI, USA*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3730567.3764781

1 Introduction

Research and Education (R&E) networks, such as Internet2 and GÉANT, provide capacity to their *members* that is not available in *commodity* (commercial) networks, enabling data intensive research such as high energy physics. R&E networks often focus on providing R&E transit *between* members, which arrange their own commodity transit. Members can therefore have multiple routes for R&E prefixes, and should generally prefer R&E routes over commodity routes.

The BGP mechanism that operators use to enforce this policy choice is known as *Local Preference* (localpref) – it is typically the first attribute that a BGP router considers when determining the *best path* if BGP has multiple routes to the same prefix. The localpref is an integer value – an operator sets a higher localpref to prefer one route over another. Operators can set the localpref for all routes received from a given neighbor by annotating the neighbor's BGP session with a default value that the router then assigns to all routes

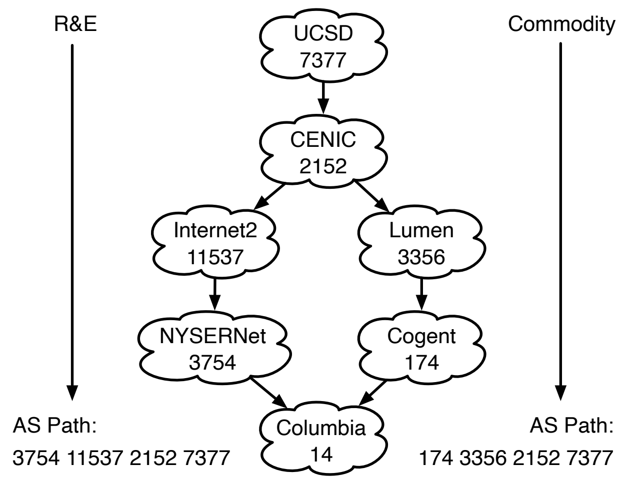

Figure 1: Columbia receives routes to the same prefix via NYSERNet (R&E) and Cogent (commodity). To deterministically prefer R&E routes, Columbia must set a higher localpref on routes from NYSERNet than on routes from Cogent.

IMC '25, Madison, WI, USA.

ACM ISBN 979-8-4007-1860-1/2025/10
https://doi.org/10.1145/3730567.3764781

received from that neighbor. If multiple routes to the same prefix have the same localpref, then BGP is most likely to use AS path length as the next tie-breaking rule, and then consider other route attributes until it selects a route.

Figure 1 illustrates the importance of using localpref. In this example, Columbia receives routes to the same UCSD prefixes via NYSERNet (R&E) and Cogent (commodity), and both routes have the same AS path length. For Columbia to deterministically select R&E routes via NYSERNet, Columbia must either (1) set a higher localpref on routes received from NYSERNet than on routes received from Cogent to be insensitive to AS path length, or (2) import only a default route from Cogent to allow R&E routes to be the most specific routes. However, in the general case, it is not possible for a third-party operator or researcher to determine relative route preference of edge networks such as Columbia, as edge networks are poorly covered by existing measurement infrastructure [1].

We develop and validate a method that infers route preference of R&E networks and use it to investigate the degree to which R&E members prefer available R&E routes. We first announced a measurement prefix to R&E and commodity providers. We then probed responsive systems in member prefixes from a host multi-homed to both the R&E and commodity providers. If responses arrived on the host's R&E interface, then the AS selected an R&E route, otherwise they selected a commodity route. We prepended the AS originating the prefix to investigate which R&E ASes used AS path length to select routes. If responses always arrived via the R&E route, then we inferred that the member (or their providers) preferred R&E routes. In May and June 2025, we used this method to survey 17,989 prefixes originated by 2,652 R&E-connected ASes, obtaining responses from systems within ≈12K (≈67%) prefixes across ≈97% R&E-connected ASes. Systems in ≈81% of the responsive prefixes always reached our host over the R&E route. We inferred that a further ≈8-9% reached our host over the R&E route if the corresponding AS path was the shortest available. With available validation data, we found at least 32 of 33 route preference inferences were correct.

While most R&E-connected ASes preferred R&E routes, the results also highlight a significant concern: some data-intensive R&E users may not benefit from the global R&E infrastructure due to local routing policies. For these users, traffic between collaborating institutions may unnecessarily traverse commodity networks, and may incur higher latency, reduced throughput, or additional cost. Given that many R&E applications – such as large-scale data transfers, remote instrumentation, or real-time collaboration – depend critically on predictable performance, the persistence of such policy-driven detours suggests that the value of the R&E infrastructure is unevenly realized across the community.

2 Background

2.1 R&E Routing

In the R&E networking space, networks such as Internet2 and GÉANT are backbone networks, providing the fabric for other regional and national R&E networks to interconnect. Members of these R&E networks are *customers* in the routing sense; for example, Internet2 exports member (customer) routes to other R&E network *peers*, such as GÉANT and AARNet. Members directly connected to Internet2 and GÉANT can have customers and peers of their own, and will generally prefer routes from their own customers and peers over routes from R&E backbone networks. R&E networks around the world can – but do not necessarily – provide global Internet routing. R&E networks can export R&E peer routes to other R&E peers – for example, Internet2 exports routes between peer NRENs to build a global R&E network.

Internet2 has four main classes of neighbor. *Participant* networks are Internet2's members, or regional networks (e.g., NYSERNet, CENIC in Figure 1) that aggregate member networks. *Peer-NREN* networks are other R&E networks, such as GÉANT, SURF, and AARNet. In this work, we study prefixes advertised to Internet2 by these two neighbor classes, where all involved traffic is R&E traffic. The other classes of Internet2 neighbor include cloud and content providers (*Peer-NET+*) and U.S. Federal agency networks (*Peer-FedNet*) who will not necessarily prefer R&E routes.

2.2 Measuring Route Preference

Decades of research [1, 11, 13, 16, 18–20, 26, 28, 31, 32, 36–39] has shown that inferring AS-and router-level routing policies is inherently difficult: operators independently tailor policies to their business needs, and BGP reveals little of the information needed for accurate modeling. The first *Gao-Rexford* AS-level models of Internet routing [12] assumed that ASes preferred routes received from customers over routes from peers and providers, encoding that policy with localpref assignments.

In 2003, Wang and Gao [36] compared AS relationship inferences with localpref assignments in looking glasses and Internet Routing Registry (IRR) records. Nearly all examined ASes (15 ASes with looking glasses, and 62 with IRR) followed the Gao-Rexford model, with > 99% of neighbor assignments for all 15 looking glass ASes and 33 of the 62 IRR ASes. In 2013, Gill *et al.* [14] surveyed 100 (mostly transit network) operators on routing policies, including localpref assignments. Most ASes followed the Gao-Rexford model, assigning uniform localpref to routes from a neighbor, with some exceptions. In 2023, Kastanakis *et al.* [18] reproduced Wang and Gao's 2003 study, finding 10 ASes with looking glasses providing localpref values, with 83% of routes conforming to the Gao-Rexford model. Some ASes assigned the same localpref to peer/provider or peer/customer routes. For 32 ASes with IRR data and ≥ 50 neighbors, 26 followed Gao-Rexford for all neighbors. Disparities between IRR and looking glass data may reflect differences between deployed and documented policies [18].

In 2006, Colitti *et al.* [6] showed that BGP path preference can be inferred by withdrawing or poisoning routes to reveal alternatives – often exposing cases where ASes preferred longer AS paths over shorter ones. In 2015, Anwar *et al.* built models of BGP decision processes of ASes informed by how those ASes reacted to (both naturally occurring and actively introduced through poisoned and varied announcement locations) changes in available routes [1]. They used publicly available BGP views, periodic traceroutes from available Vantage Points (VPs), and CAIDA AS relationship inferences [24] to build these models, reporting that 14-35% of observed decisions deviated from expectations under Gao-Rexford and AS path length-based models.

In 2021, Fonseca *et al.* [10] showed that an AS can localize spoofed traffic sources by first pre-computing how networks react to varied (e.g., prepending, poisoning, announcement locations) route announcements using RIPE Atlas VPs. In essence, relatively few networks react the same way to a series of targeted route announcements. The AS can match attack traffic to this map by adjusting route announcements for an attacked prefix according to the map.

2.3 Active Measurement of Routing

Previous work inferring route preference (§2.2) used available VPs – public BGP data, looking glasses, and active measurement platforms – limiting their applicability, particularly for edge networks [1]. However, there is a body of work that infers other properties of BGP routing using responsive destinations in edge networks. For example, Labovitz *et al.* [21] studied the effects of BGP convergence on the data-plane, Bush *et al.* [2] and Rodday *et al.* [29] studied the use of default routes, Bush *et al.* [2] studied the presence of hidden upstreams and stale bogon filters, De Vries *et al.* [7] studied anycast catchments, and Cartwright-Cox [4] studied the effects of RPKI Route Origin Validation (ROV) deployment. These works send ICMP pings to responsive addresses, sourced from addresses within BGP routes with properties whose effect they seek to measure, and trade the ability to reason about root causes for increased coverage – tens of thousands of ASes. Each responsive address becomes what De Vries *et al.* [7] call a *passive VP*.

The use of ICMP pings to study ROV has attracted some criticism, as a passive VP unresponsive to pings sourced from an RPKI-invalid prefix may not necessarily be because the AS hosting the passive VP has deployed ROV itself – an AS further along the return path may have [30]. The same AS can appear as the beneficiary of ROV when probed from some locations, but not others. Cartwright-Cox [4] also acknowledged that default routes may allow traffic to reach an RPKI-invalid prefix even when ASes have deployed ROV.

We use the passive VP concept in this work to infer R&E routing policy. We are not concerned with underlying causes. Rather, we seek to understand whether scientific flows remain on scientific infrastructure.

3 Inferring R&E Route Preference

3.1 Measurement Host

We probed systems in R&E prefixes from a host physically located in Atlanta, GA and managed by Internet2, configured as illustrated in Figure 2. We assigned a publicly-routed IPv4 address within a measurement prefix to the host's loopback interface, and used that source address in our probes. We verified that commodity providers did not learn the R&E path by not announcing the commodity prefix before our experiments began; in the available public BGP data, only R&E networks reported a path to the measurement prefix, and none reported a commodity ASN in the AS path.

The host was connected to one router via multiple VLAN interfaces. The host supported experiments from two R&E networks – SURF (a Dutch national R&E network) on 30th May 2025, and Internet2 (a U.S. R&E backbone network) on 5th June 2025. We always announced a commodity route for the measurement prefix. SURF announced the R&E route for the May experiment, and Internet2 announced the R&E route for the June experiment. SURF operators provisioned a tunnel to deliver responses to our measurement host at Internet2 while they originated the R&E route; Internet2 presented this tunnel to our host via a VLAN interface. Internet2 used separate virtual routing and forwarding (VRF) instances for commodity and R&E routing, which presented to our host as additional VLAN interfaces. The VLAN on which a response arrived identifies whether the response followed R&E or commodity routing. We extended scamper [22] to record the interface on which the operating system received responses using the IP_PKTINFO ancillary message [33] via recvmsg. We developed a program that used the scamper Python module [23] to conduct the measurement and produce JSON results, which is publicly available [25].

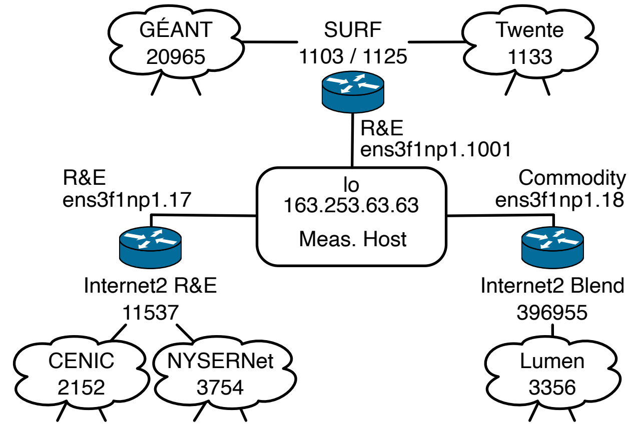

Figure 2: Measurement host configuration. The host was connected to three VLANs which interfaced with R&E and commodity networks. We inferred the type of route R&E networks used to reach our measurement host based on the VLAN interfaces that the responses arrived on.

3.2 Probe Seeds

We used the results of two comprehensive, complementary scanning projects to seed our probing. The first, the ISI Internet Addresses IPv4 Response History Dataset [34], summarizes previous ISI censuses of the IPv4 address space [17], ranking addresses that were ever responsive in any ISI census in order of those most likely to currently respond [9]. The second, Censys [8], provides vetted researchers with access to a search engine of Internet-wide scanning results for approved non-commercial use [5], which we used to query for responsive TCP and UDP services in each R&E prefix.

We began with 18,427 Participant and Peer-NREN prefixes (§2.1) propagated by Internet2 to RouteViews on May 29th 2025 at 10:00 UTC, where all involved prefixes were R&E. We excluded 437 prefixes entirely covered by other prefixes, and the measurement prefix, leaving 17,989 prefixes. We then extracted seeds from the ISI history dataset, and queried the Censys API about hosts in these prefixes over the course of 7 hours. Of the 17,989 prefixes, 11,731 (65.2%) had a seed in the ISI history dataset, covering 2,542 (95.8%) of 2,653 ASes. When we included Censys, coverage increased to 13,189 (73.3%) prefixes originated by 2,622 (98.8%) ASes.

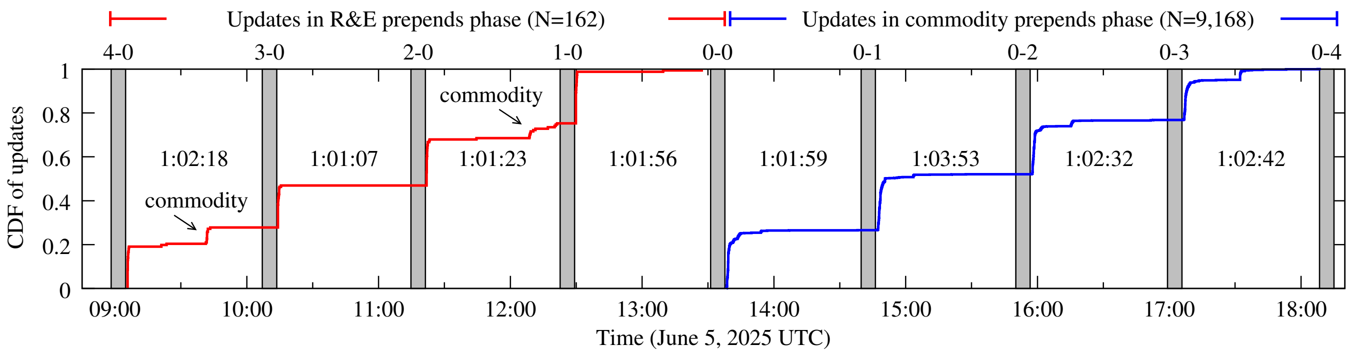

Figure 3: Flow between active probing windows (≈7 minutes, grey bars) and measurement prefix BGP activity during our June 2025 (Internet2) experiment. In the first five (R&E prepends) rounds, updates after R&E convergence exclusively occurred on commodity routes, and were 26 of the 162 updates during the R&E prepends phase.

Beginning May 29th at 21:30 UTC, for each R&E prefix, we probed up to ten addresses from the ISI history file ranked by their score, and up to ten randomly selected address-port tuples in Censys data, in order to have a current set of known-responsive addresses in each prefix. Our goal was to find three responsive addresses within each prefix to reduce the chance that we were unlucky and only selected an address in the prefix assigned to a router operated by a different AS than the AS routing the prefix [2]. As not all addresses were responsive when we probed them (some prefixes covered by addresses in the ISI history file were last responsive more than a year ago), we ended up with addresses in 12,241 (68.0%) prefixes originated by 2,594 (97.8%) ASes. We obtained three destinations in 10,123 (82.7%) of 12,241 responsive prefixes. For 10,253 (77.8%) prefixes, we used ICMP seeds from the ISI history data, and for 3,215 (24.4%) we used TCP and UDP seeds from the Censys data – 283 (2.1%) prefixes had mixed seed origin. We used the same probe seeds for both the SURF (30 May 2025) and Internet2 (5 June 2025) experiments to allow for comparable results.

3.3 BGP Advertisements

We advertised the measurement prefix with origin AS 396955 to commodity networks via Lumen (AS 3356), with origin AS 1125 to R&E networks for the experiment via SURF (AS 1103), and with origin AS 11537 to R&E networks for the experiment with Internet2. These announcements were covered by RPKI ROAs and IRR route objects. For both SURF and Internet2, we conducted a series of nine tests, using a different AS prepend configuration in each, to infer member sensitivity to AS path length. We use "4-0" to refer to the test with 4 prepends of the R&E ASN and no prepending of the commodity ASN, "0-0" to refer to no prepending of either R&E or commodity ASN, and "0-4" to refer to the test with 4 prepends of the commodity ASN and no prepending of the R&E ASN. The order of our nine tests was "4-0", "3-0", "2-0", "1-0", "0-0", "0-1", "0-2", "0-3", and "0-4" – we decreased prepends of the R&E ASN, and then increased prepends of the commodity ASN, to minimize the variables that could affect routing decisions between tests.

We sought to minimize the effect of route flap damping (RFD), where routers maintain a penalty per prefix/BGP session pair, as RFD could lead to our BGP configurations being suppressed [3, 35]. To minimize the effect of RFD but still allow our experiment to complete within a work day (operators at Internet2 and SURF manually adjusted each route announcement) we conducted active probing one hour after changing BGP configurations. In 2020, Gray reported that ≈9% of the ASes they measured enabled RFD, few ASes damped prefixes longer than 15 minutes, and that they did not observe suppress times longer than one hour [15].

Figure 3 illustrates the timeline of our June 5th 2025 (Internet2) experiment, which began shortly before 9:00 UTC with the prepend configuration at "4-0" (§3.1) for an hour prior to our experiment. We plot cumulative BGP update activity observed by all RouteViews and RIPE RIS peers for the measurement prefix during the experiment. We modified prepend configurations immediately after each round of active probing completed, and then waited an hour before starting the next round. We plot cumulative BGP update churn while varying prepends on the R&E route (configurations "4-0" to "0-0"), and then while varying prepends on the commodity route (configurations "0-0" to "0-4"). There was comparatively sparse BGP activity for prepend changes on the R&E route, as few public BGP views observed the R&E route. In BGP data recorded by RouteViews and RIPE RIS, we observed 162 update messages across more than four hours while varying prepends on the R&E route, and 9,162 update messages across the four hours while varying prepends on the commodity route. Figure 3 shows that (at least for the public view) BGP update activity for the measurement prefix was relatively settled for at least 50 minutes prior to the active measurement for that configuration. In the first five (R&E prepends) tests, updates after R&E convergence exclusively occurred on commodity routes, and were 26 of the 162 updates during that phase. Each round of active probing took ≈7 minutes at 100pps.

3.4 Limitations

Our goal is to infer routing policy affecting return paths from R&E-connected systems. Our approach assumes that our measurement prefix is representative of how the origin considers other routes from the BGP session they learn our routes from. However, the network hosting the remote system (or an intermediate transit) may apply localpref on finer granularities than per BGP-session [14]. While packet probing allows for route preference inferences at a finer granularity than AS-level, it requires responsive hosts deployed across the AS to capture that diversity.

Inference	Prefixes		ASes	
Always R&E	9,852	81.8%	1,958	76.1%
Always commodity	843	7.0%	339	13.2%
Switch to R&E	963	8.0%	302	11.7%
Switch to commodity	1	0.0%	1	0.1%
Mixed R&E + commodity	382	3.1%	234	9.1%
Oscillating	6	0.0%	5	0.2%
Total:	12,047		2,574	

(a) Results for SURF Experiment (29 May 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 + commodity	371	3.1%	228	8.8%
Oscillating	2	0.0%	2	0.1%
Total:	12,077		2,578	

(b) Results for Internet2 Experiment (5 June 2025)

Table 1: Results for tested prefixes. Systems in most (≈88%) prefixes were insensitive to our AS path length changes, and most (≈81%) always used the R&E route. The percentages in the ASes column add to more than 100% because some ASes are included in multiple categories.

Packet loss		279	
Mixed R&E + commodity		400	
Oscillating		6	
Switch to commodity		4	
Incomparable prefixes:		689	
SURF (May '25)	**Internet2 (June '25)**		
Always commodity	Always R&E	27	0.2%
Always commodity	Switch to R&E	37	0.3%
Always R&E	Always commodity	19	0.2%
Always R&E	Switch to R&E	184	1.6%
Switch to R&E	Always commodity	35	0.3%
Switch to R&E	Always R&E	61	0.5%
Different inferences:		363	3.1%
Always commodity	Always commodity	761	6.6%
Always R&E	Always R&E	9,569	82.8%
Switch to R&E	Switch to R&E	859	7.4%
Same inferences:		11,189	96.9%
Comparable prefixes:		11,552	

Table 2: Comparison of SURF and Internet2 results. For comparable prefixes, 96.9% of inferences were the same. 161 of 363 (44.3%) differences were because NIKS assigned a higher localpref to GÉANT (where NIKS learned the SURF route) than NORDUnet (where NIKS learned the Internet2 route), and assigned NORDUnet and commodity the same localpref.

4 Do R&E Networks Prefer R&E Routes?

Table 1 summarizes results for both SURF and Internet2 experiments. We characterized prefixes that had a response from at least one system during every active probing round, so these tables exclude ≈160 of 12,241 prefixes for which we had seeds. In both experiments, systems in most prefixes were insensitive to our BGP configuration changes – response traffic for ≈81% of prefixes always arrived via the R&E route, and always via commodity for ≈7% of prefixes. Table 1 also summarizes our results by origin AS. In both experiments, ≈75-76% of tested ASes originated at least one prefix whose systems always replied over R&E.

The ordering of our prepend adjustments (§3.3) implies that networks that assign equal localpref will send responses over commodity, and then over R&E, and not make further transitions between route types, provided those networks consider AS path length. Systems in ≈8-9% of the prefixes switched from commodity to R&E when we used a prepend configuration that caused a router on the return path to prefer the R&E route – implying that BGP used AS path length to select the route because available R&E and commodity routes had the same localpref.

Four ASes unexpectedly switched from R&E to commodity during our two experiments, even as we increased commodity prepends. When we observed this behavior in preliminary experiments, we discussed our findings with an operator at an AS, who reported that an outage during our experiment caused their route to our host to revert to commodity. We therefore inferred that a network assigned equal localpref to commodity and R&E routes only when we received responses over commodity, and then over R&E.

During individual active probing rounds, responses from some systems in some prefixes returned over R&E, while responses from other systems in those same prefixes returned over commodity. *Across* active probing rounds, responses from systems in some prefixes also appeared to be affected by AS path length as we decreased R&E prepends and increased commodity prepends, as responses from these systems would switch to R&E while responses from other systems continued to return over either R&E or commodity. We labelled any prefix that had at least one active probing round with different types of return path as a *mixed* prefix, accounting for 3.1% of responsive prefixes in both experiments. Because the absolute number of prefixes with mixed return paths changed *across* active probing rounds as we modified BGP prepend configurations, we summarize the overall pattern that we observed. Overall, two systems in each mixed prefix would prefer R&E, while one would prefer commodity, at an overall ratio of ≈2:1. This is consistent with R&E members generally preferring R&E routes.

In both experiments, 2-6 (0.0%) prefixes had oscillating responses, where there were multiple state transitions between R&E and commodity during the experiment. Possible explanations for this behavior are short term outages during the experiments, and the use of route optimization appliances that use route properties determined outside of BGP to determine best paths on fine-grained timescales.

Table 2 compares prefix-level inferences between the two experiments, which were run one week apart with the same probe seeds. Different prefixes experienced packet loss in the two experiments,

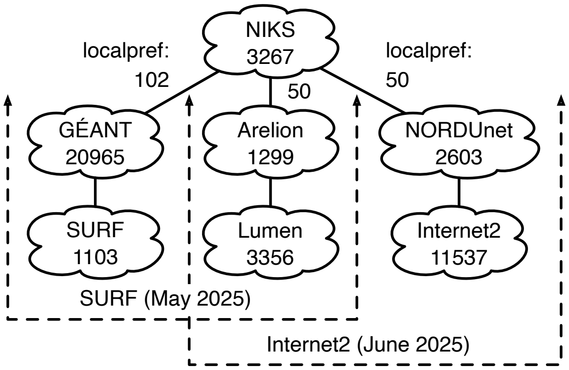

Figure 4: Effect of NIKS assigning different localpref values to different neighbors. Responses always arrived over R&E for the SURF experiment because NIKS assigned GÉANT a higher localpref. Responses arrived over R&E or Commodity depending on AS path length during Internet2 experiment.

Inference	Congruent	Incongruent	Total
Always R&E	15	3*	18
Always commodity	3	0	3
Switch to R&E	4	0	4
Total	22	3*	25

Table 3: Comparing policy inferences with public BGP views. 22 of 25 ASes had policy inferences congruent with the BGP view. At least two of the three incongruent public views did not reflect the actual routing policy of the AS.

and prefixes with mixed R&E and commodity routing have ambiguous routing policies, so we do not consider these 689 prefixes to be comparable prefixes. For 11,552 comparable prefixes, 96.9% of the inferences were the same in both experiments.

Nearly half of the prefix-level differences (161 of 363) and 37 of all 125 ASes with any prefix-level difference were due to configuration by NIKS, a Russian R&E transit network. Shown in Figure 4, NIKS assigned a higher localpref to GÉANT (where NIKS learned the SURF route) than NORDUnet (a Nordic R&E transit network where NIKS learned the Internet2 route), while assigning NORDUnet and commodity routes the same localpref [27]. NIKS used NORDUnet as one of its global transit providers, so assigned the same localpref as it did to routes from Arelion so that NIKS used both connections for outgoing traffic. During the SURF experiment, NIKS always used the R&E route via GÉANT, but used the commodity route via Arelion during the Internet2 experiment when Arelion's AS path was shorter than NORDUnet's. This policy affected 161 of 184 the prefixes in the associated row in Table 2. Some of the other differences can be explained by peering between R&E networks that bypasses other R&E transit, where operators do not assign that other R&E transit a higher localpref than commodity routes available to that AS.

§A discusses the interplay of our prepend ordering with other BGP tie-breaks when ASes assign the same localpref to R&E and commodity routes. We chose the prepend configuration sequence to minimize variables changing between experiments. §B discusses sensitivity to AS path lengths. During the SURF experiment, *Participant* (U.S. domestic R&E) ASes switched to R&E one prepend configuration later than *Peer-NREN* (international) ASes because their R&E AS paths were longer as a population.

4.1 Validation

4.1.1 Comparison with public BGP view. Of the 2,578 ASes with at least one responsive prefix, 26 ASes also provided a public (RouteViews or RIPE RIS) BGP view, which we compared to our prefix-level inferences. For each collector, we downloaded the June 5th 08:00 UTC RIB file and all update files through the entirety of our Internet2 experiment, extracting activity for our measurement prefix. Most tested ASes originated multiple prefixes; if an AS had different inferences for different prefixes, we considered the most frequent prefix-level inference for each AS. One AS had no most frequent inference and we did not include them in this validation. For the remaining 25, we say that the prefix-level inferences are congruent with the public BGP view if the origin AS or ASes observed in BGP would be expected given the inference. For example, we expected to only observe routes for our measurement prefix from the AS's BGP view originated by AS 11537 when we observed response packets only arrive via R&E.

Table 3 shows that 22 of 25 ASes had BGP activity congruent with our inferences. For three ASes, we inferred their policy was to always prefer the R&E route, but we saw only the commodity route in the public BGP view. We contacted operators at the three incongruent ASes and received a response from two. Both reported that the AS used multiple VRFs – one for R&E routing, and one for commodity routing. While their policy was to prefer R&E routes, they exported routes from the commodity VRF to the public BGP collector. That is, our policy inference was correct.

4.1.2 Operator Ground Truth. We contacted operators at ten ASes with our findings, and received responses from eight, covering the spectrum of our inferences. For two ASes, we inferred that they assigned equal localpref to both R&E and commodity routes, as we inferred their return route was sensitive to AS path length, which they confirmed. For another AS, we observed one /24 prefix where two systems replied via R&E, while a third system returned via commodity. They confirmed that systems in prefixes originated by their AS used the R&E route, but that we had probed a router that used an address from their prefix for interconnection, and that the router did not have an R&E route. Finally, we heard from five ASes that our inferences that they preferred R&E routes over commodity routes was correct.

4.2 Does Preference Align with Prepending?

When we discussed our inferences of egress routing policies with operators at R&E ASes, some operators volunteered their prepending policies, which can influence egress routing decisions *made by other ASes* when those other ASes assign equal localpref to R&E and commodity routes and tie-break using AS path length. A natural behavior for an AS X that prefers R&E for their egress routing is to prepend their commodity route announcements, so that other R&E ASes that do not assign a higher localpref to R&E routes might still use an R&E path to X as a result of the commodity route announcement having a longer AS path.

Inference	Prepends			No
	R = C	R < C	R > C	commodity
Always R&E	3,005 73.8%	2,628 83.2%	204 50.7%	3,921 88.3%
Always commodity	319 7.8%	192 6.1%	149 37.1%	180 4.1%
Switch to R&E	610 15.0%	248 7.9%	28 7.0%	217 4.9%
Mixed R&E + commodity	138 3.4%	90 2.8%	21 5.2%	122 2.7%
Total	4,072	3,158	402	4,440

Table 4: Fraction of prefixes with origin AS prepending in AS paths toward R&E and commodity, by route preference inference. R < C means that the origin was prepended more in commodity routes than R&E routes.

Table 4 compares inferences of egress routing policies with origin AS prepending in corresponding BGP routes for R&E prefixes as recorded by RouteViews and RIPE RIS in the June 5th 08:00 UTC RIB files, corresponding to the Internet2 (June 2025) experiment. For this analysis, we first classified ASes in any *Participant* or *Peer-NREN* (§2.1) route observed by Internet2 as R&E ASes – i.e., the set of R&E members and R&E transit providers. Then, if the immediate upstream AS was not in the set of R&E ASes, we classified the route as via a commodity AS. The "no commodity" column covers the 4,440 prefixes whose only BGP-observed upstream was initially via an R&E transit provider. We do not further consider the prepending properties of these routes, as there is no commodity route to compare to. The majority (88.3%) of responses for systems in these prefixes always arrived via R&E networks, but 9.0% did not. Some networks in this 9.0% may have unobserved commodity transit providers which, if undetected, would impact the accuracy of routing models.

Of the remaining 7,632 prefixes, 3,560 (46.6%) had the origin AS unequally prepended between R&E and commodity neighbors, perhaps due to operator desire to influence egress routing of other ASes. 3,158 (91.5%) of prefixes had their origin AS prepended more towards commodity neighbors than to R&E. Systems in most (83.2%) of these 3,158 prefixes always returned via R&E, though 7.9% could be influenced by the prepend configuration of our measurement prefix. For the 402 prefixes where the origin AS was prepended more via R&E than via commodity, 37.1% contained systems that always returned via commodity, suggesting a deliberate policy choice by those ASes to use commodity routing as much as possible. However, 50.7% of these 402 prefixes contained systems whose return route was always via R&E. Where origin AS prepending was equal between R&E and commodity routes, 73.8% always returned via R&E. We believe that our results show that while relative AS prepending provides some signal as to the relative egress route preference of these networks, relying on that signal would lead to error in route predictions of these networks.

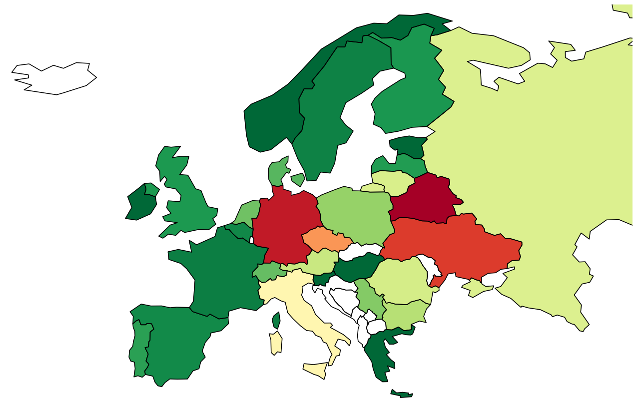

(a) Europe.

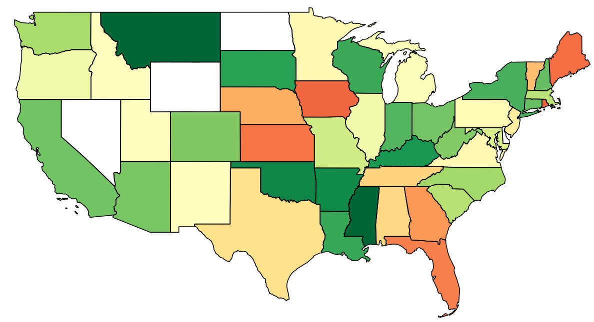

(b) US states.

Figure 5: Percentage of ASes in a given region reached by RIPE over an R&E route for at least one prefix on 29 May 2025. The color scale ranges from dark red (0%) to dark green (100%). Regions that are a shade of red are mostly reached over commodity because their commodity announcements were not prepended.

4.3 Equal localpref and R&E Route Selection

We examined the properties of routes selected by RIPE, an R&E-connected European supplier of Internet measurement data heavily used in the research community. These are the routes RIPE uses to deliver public BGP data to R&E-connected institutions. We inferred that RIPE assigned equal localpref to commodity and R&E routes, and validated that inference with them. RIPE provides a public BGP view, which we used to examine how they reached R&E networks in practice on 29 May 2025.

RIPE had matching routes for 18,160 of 18,427 R&E prefixes (§3.2) originated by 2,640 ASes. We identified whether RIPE used an R&E or commodity path by manually classifying their neighbors as R&E or commodity ASes. Overall, RIPE used R&E routes to reach 11,616 (64.0%) prefixes originated by 1,688 ASes (63.9% of the 2,640 ASes), and commodity routes to reach the remaining 6,544 prefixes originated by 1,165 ASes (44.1%, some ASes originated multiple prefixes that RIPE used a mixture of paths to reach). We used the Netacuity Edge geolocation database of 30 May 2025 to map R&E prefixes to countries and U.S. states.

We calculated the percentage of R&E-connected ASes that had at least one prefix reached over an R&E per country (Figure 5a, restricted to Europe for visibility) and per U.S. state (Figure 5b),

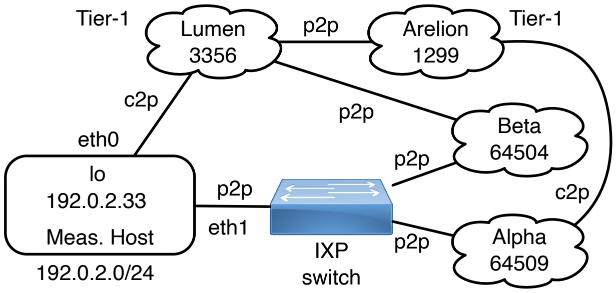

Figure 6: Measurement host configuration for inferring relative route preference of peer and provider routes.

where the country or state had at least four geolocated R&E ASes. Some countries (e.g., Norway, Sweden, France, Spain, Australia, New Zealand) had more than 90% of their ASes reached over R&E due to (1) national R&E networks (NRENs) also providing commodity transit, (2) customers of those NRENs near exclusively using them for transit, and (3) those NRENs prepending their announcements to commodity transit providers so that other networks would likely select the R&E route. There were also countries (e.g., Brazil, Thailand, Germany, Ukraine, Belarus) where fewer than 15% of their ASes were reached over R&E because the R&E paths lost BGP tie breakers. For German R&E, this was due to RIPE and DFN (the German NREN) using Deutsche Telekom (DT, AS 3320) as a common provider. DFN did not prepend the route they announced to DT. The same situation repeats for the other listed countries.

In the U.S., the New York state R&E network (NYSERNet) does not offer commodity transit, yet 84% of 74 New York-mapped ASes were reached over an R&E route. This is because their members are conditioned to prepend their own AS in commodity announcements. Interestingly, a lower percentage (78%) of 127 California-mapped ASes were reached over R&E, despite the California state R&E network (CENIC) offering commodity transit and CENIC prepending their announcements to commodity. This is due to some California-mapped ASes arranging additional commodity transit and not prepending their AS on their commodity announcements.

5 Discussion

R&E institutions often need to distribute large volumes of scientific data globally – a task that is impractical using commercial content distribution models. Instead, they depend on specialized R&E networking infrastructure to bridge that gap. We demonstrated and validated an approach to infer relative route preference in the R&E ecosystem, which covers 2,653 (≈3.5%) of routed ASes. We also examined how AS path prepending affects route selection in this environment. Our findings have implications for routing policy design in R&E networks, and to support future research, we publicly release both our source code and dataset [25].

We believe our method has broader applicability beyond our demonstration. It could be used to detect whether ASes assign equal localpref values to both provider and peer routes if those ASes rely on AS path length as the next tie-break. A practical setup for this would involve connecting a host to both a large IXP and a Tier-1 transit provider that peers selectively (Figure 6). This configuration allows testing the route preference policies of most ASes directly connected to the IXP, such as Alpha in the figure. Because the host would have separate physical interfaces for the IXP and transit connections, it could identify the class of return route by observing which interface the response traffic arrives on. This method is practical so long as the tested ASes (such as Beta) do not also peer with the measurement host's transit provider. If they did, the probed AS would have two peer routes to choose from, making it impossible to isolate the preference between peer and provider paths. One approach would be to use a second Tier-1 AS as a second provider that announces the provider route separately to the first, in the hope Beta is not a peer of both Tier-1s. Furthermore, our method could capture subtler behaviors – such as when different IP prefixes from the same IXP member AS are routed differently, possibly due to geographic distribution or internal traffic engineering, enriching empirical routing models.

Acknowledgments

We thank the operators who candidly discussed aspects of their network's operations, and the anonymous reviewers for their feedback. This research was supported by National Science Foundation (NSF) grants OAC-2131987 and OAC-2530871, but this paper represents only the position of the authors.

References

[1] Ruwaifa Anwar, Haseeb Niaz, David Choffnes, Ítalo Cunha, Phillipa Gill, and Ethan Katz-Bassett. 2015. Investigating Interdomain Routing Policies in the Wild. In *IMC*. 71–77.

[2] Randy Bush, Olaf Maennel, Matthew Roughan, and Steve Uhlig. 2009. Internet Optometry: Assessing the Broken Glasses in Internet Reachability. In *IMC*. 242–253.

[3] Randy Bush, Cristel Pelsser, Mirjam Kuhne, Olaf Maennel, Pradosh Mohapatra, Keyur Patel, and Rob Evans. 2013. RIPE Routing Working Group Recommendations on Route Flap Damping. https://www.ripe.net/publications/docs/ripe-580/.

[4] Ben Cartwright-Cox. 2018. Are BGPs security features working yet? https://blog.benjojo.co.uk/post/are-bgps-security-features-working-yet-rpki.

[5] censys.io. 2024. Research Access to Censys Data. https://support.censys.io/hc/en-us/articles/360038761891-Research-Access-to-Censys-Data.

[6] Lorenzo Colitti, Giuseppe Di Battista, Maurizio Patrignani, Maurizio Pizzonia, and Massimo Rimondini. 2006. Investigating Prefix Propagation through Active BGP Probing. In *ISCC*. 497–504.

[7] Wouter B. de Vries, Ricardo de O. Schmidt, Wes Hardaker, John Heidemann, Pieter-Tjerk de Boer, and Aiko Pras. 2017. Broad and Load-Aware Anycast Mapping with Verfploeter. In *IMC*. 477–488.

[8] Zakir Durumeric, David Adrian, Ariana Mirian, Michael Bailey, and J. Alex Halderman. 2015. A Search Engine Backed by Internet-Wide Scanning. In *CCS*. 542–553.

[9] Xun Fan and John Heidemann. 2010. Selecting Representative IP Addresses for Internet Topology Studies. In *IMC*. 411–423.

[10] Osvaldo Fonseca, Ítalo Cunha, Elverton Fazzion, Wagner Meira Jr., Brivaldo Junior, Ronaldo A. Ferreira, and Ethan Katz-Bassett. 2021. Identifying Networks Vulnerable to IP Spoofing. *IEEE TNSM* 18, 3 (Sept. 2021), 3170–3183.

[11] Lixin Gao. 2001. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking* 9, 6 (Dec. 2001), 733–745.

[12] Lixin Gao and Jennifer Rexford. 2001. Stable Internet Routing Without Global Coordination. *IEEE/ACM Transactions on Networking* 9, 6 (Dec. 2001), 681–692.

[13] Phillipa Gill, Michael Schapira, and Sharon Goldberg. 2012. Modelling on quicksand: Dealing with the scarcity of ground truth in interdomain routing data. *CCR* 42, 1 (Jan. 2012), 40–46.

[14] Phillipa Gill, Michael Schapira, and Sharon Goldberg. 2014. A Survey of Interdomain Routing Policies. *CCR* 44, 1 (Jan. 2014), 28–34.

[15] Caitlin Gray, Clemens Mosig, Randy Bush, Cristel Pelsser, Matthew Roughan, Thomas C. Schmidt, and Matthias Wahlisch. 2020. BGP Beacons, Network Tomography, and Bayesian Computation to Locate Route Flap Damping. In *IMC*. 492–505.

[16] Sichang He, Ítalo Cunha, and Ethan Katz-Bassett. 2024. RPSLyzer: Characterization and Verification of Policies in Internet Routing Registries. In *IMC*. 365–374.

[17] John Heidemann, Yuri Pradkin, Ramesh Govindan, Christos Papadopoulos, Genevieve Bartlett, and Joseph Bannister. 2008. Census and Survey of the Visible

Internet. In *IMC*. 169–182.
[18] Savvas Kastanakis, Vasileios Giotsas, Ioana Livadariu, and Neeraj Suri. 2023. 20 Years of Inferring Inter-domain Routing Policies. In *IMC*.
[19] Savvas Kastanakis, Vasileios Giotsas, and Neeraj Suri. 2022. Understanding the confounding factors of inter-domain routing modeling. In *IMC*. 758–759.
[20] Ege Cem Kirci, Valerio Torsiello, and Laurent Vanbever. 2024. What is the next hop to more granular routing models?. In *HOTNETS*. 343–351.
[21] Craig Labovitz, Abha Ahuja, Abhijit Bose, and Farnam Jahanian. 2000. Delayed Internet Routing Convergence. In *SIGCOMM*. 175–187.
[22] Matthew Luckie. 2010. Scamper: a Scalable and Extensible Packet Prober for Active Measurement of the Internet. In *IMC*. 239–245.
[23] Matthew Luckie, Shivani Hariprasad, Raffaele Sommese, Brendon Jones, Ken Keys, Ricky Mok, and K Claffy. 2025. An Integrated Active Measurement Programming Environment. In *PAM*. 137–152.
[24] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, and k claffy. 2013. AS Relationships, Customer Cones, and Validation. In *IMC*. 243–256.
[25] Matthew Luckie, Steven Wallace, Karl Newell, Jeff Bartig, Sadi Koçak, Niels den Otter, Kaj Koole, James Deaton, and k claffy. 2025. Supplement for "R&E Routing Policy: Inference and Implication". https://publicdata.caida.org/datasets/supplement/2025-imc-re-routing-policy/.
[26] Wolfgang Mühlbauer, Anja Feldmann, Olaf Maennel, Matthew Roughan, and Steve Uhlig. 2006. Building an AS-topology model that captures route diversity. In *SIGCOMM*. 195–206.
[27] NIKS. 2025. AS3267 Looking Glass. https://lg.niks.su/.
[28] Bruno Quoitin and Steve Uhlig. 2005. Modeling the routing of an autonomous system with C-BGP. *IEEE Network* 19, 6 (Nov. 2005), 12–19.
[29] Nils Rodday, Lukas Kaltenbach, Ítalo Cunha, Randy Bush, Ethan Katz-Bassett, Gabi Dreo Rodosek, Thomas C. Schmidt, and Matthias Wählisch. 2021. On the Deployment of Default Routes in Inter-domain Routing. In *TAURIN*. 14–20.
[30] Nils Rodday, Ítalo Cunha, Randy Bush, Ethan Katz-Bassett, Gabi Dreo Rodosek, Thomas C. Schmidt, and Matthias Wählisch. 2021. Revisiting RPKI Route Origin Validation on the Data Plane. In *TMA*.
[31] Pavlos Sermpezis and Vasileios Kotronis. 2019. Inferring catchment in Internet routing. In *SIGMETRICS*.
[32] Rachee Singh, David Tench, Phillipa Gill, and Andrew McGregor. 2021. PredictRoute: a network path prediction toolkit. In *SIGMETRICS*.
[33] W. Stevens, M. Thomas, E. Nordmark, and T. Jinmei. 2003. Advanced Sockets Application Program Interface (API) for IPv6. https://www.rfc-editor.org/rfc/rfc3542.
[34] USC/LANDER project. 2025. Internet Addresses IPv4 Response History Dataset, PREDICT ID: USC-LANDER/internet_address_history_it111w-20250226/rev15218.
[35] Curtis Villamizar, Ravi Chandra, and Ramesh Govindan. 1998. BGP Route Flap Damping. https://www.rfc-editor.org/rfc/rfc2439.
[36] Feng Wang and Lixin Gao. 2003. On inferring and characterizing Internet routing policies. In *IMC*. 15–26.
[37] Tianhao Wu, Jessie Hui Wang, Jilong Wang, and Shuying Zhuang. 2022. RouteInfer: Inferring Interdomain Paths by Capturing ISP Routing Behavior Diversity and Generality. In *PAM*. 216–244.
[38] Jiangchen Zhu, Kevin Vermeulen, Ítalo Cunha, Ethan Katz-Bassett, and Matt Calder. 2002. The Best of Both Worlds: High Availability CDN Routing Without Compromising Control. In *IMC*. 655–662.
[39] Ítalo Cunha, Pietro Marchetta, Matt Calder, Yi-Ching Chiu, Bruno VA Machado, Antonio Pescapè, Vasileios Giotsas, Harsha V Madhyastha, and Ethan Katz-Bassett. 2016. Sibyl: a practical Internet route oracle. In *NSDI*. 325–344.

A AS Path Prepend Ordering

We had some influence on two route attributes possibly considered by BGP in selecting a best route – the AS path length, and route age. BGP considers these attributes only if available routes have equal localpref. BGP considers AS path length immediately after localpref, and considers route age only if multiple routes have equal intradomain costs (which we had no influence on). When considering route age, BGP selects the oldest route so that it prefers stable routes.

Figure 7 visualizes the interplay between route age and AS path lengths for our two experiments, relative to the AS path lengths of R&E and commodity routes available to an R&E institution. During the first phase where we decreased R&E prepends, the commodity route would have been older than the R&E route, because we did not change prepends on the commodity route. Networks that would

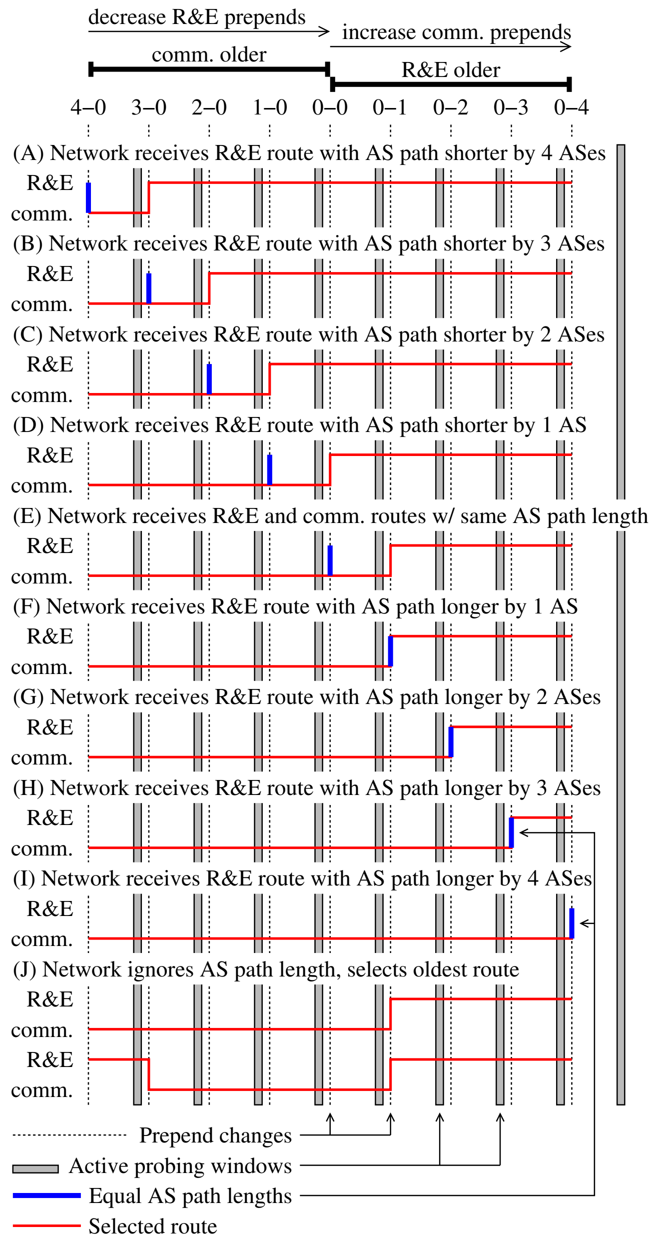

Figure 7: State diagrams illustrating the influence of AS path length and route age on route selection. All other things equal, a network receiving R&E routes shorter than commodity routes selects an R&E route when commodity routes become longer, if the network breaks ties with route age.

have received shorter R&E routes in the absence of our prepending would not have switched to the R&E route until the commodity route's AS path length was longer (cases A – E). During the phase where we increased commodity prepends, the R&E route would have been older than the commodity route, so networks that would receive shorter commodity routes in the absence of our prepending would have immediately switched to the R&E route because the R&E route was older (cases F – I). We also show state diagrams for

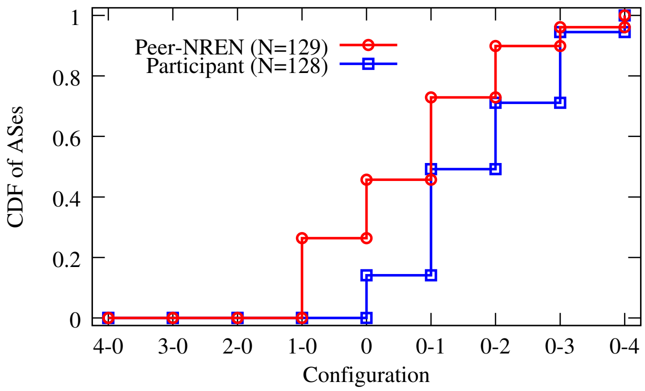

(a) Results for SURF Experiment (29 May 2025)

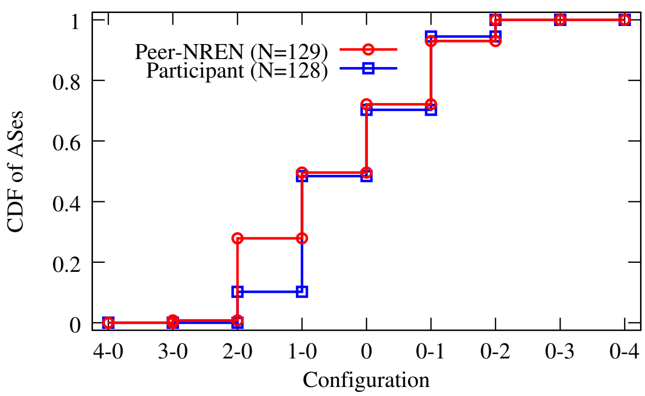

(b) Results for Internet2 Experiment (5 June 2025)

Figure 8: Effect of AS path configuration on selection of R&E routes, as a CDF. If populations experienced similar AS path lengths in each experiment, their lines would overlap.

networks that may have ignored AS path length, and have equal intradomain tie breakers. If the commodity route was older when we began each experiment, the network would have selected the commodity route until we began increasing commodity prepends (the first row in case J). However, if the R&E route was older when we began each experiment, the network will have switched to the commodity route when we decreased its AS path length, and then switched back to the R&E route when we increased the AS path length of the commodity route (the second row in case J).

Regular BGP churn that caused the route age to reset to zero is invisible to us as experimenters, and outside of our control. This churn matters when the available commodity and R&E routes had equal AS path length, a network used route age as a tie breaker, and the churn occurs in the hour between us changing the prepends on the route and probing responsive destinations. For cases A – E in Figure 7, if some event occurred on the commodity route, then the network will switch to the R&E route and we will observe responses return to our measurement host on an R&E path in the active probing window before we next change route announcements. For cases F – I in Figure 7, if some event occurred on the R&E route, then the network will return to the commodity route until the commodity prepends are increased further.

B When Did ASes Switch to R&E Routes?

We now investigate the effects of relative AS path length on the return paths selected by ASes. To allow experiments to be compared, we selected the 859 prefixes that switched from commodity to R&E in both (SURF and Internet2) experiments, originated by 254 ASes. Then, we considered the first configuration for each AS where a prefix switched to R&E, so that we count ASes originating many prefixes that switch in unison once. Further, we considered prefixes in two broad classes identified in §2.1: U.S. domestic R&E networks (the *Participant* class, originated by 128 ASes), and international networks (the *Peer-NREN* class, originated by 129 ASes) – three ASes originated prefixes in both classes.

Figure 8 shows the cumulative distribution of when each AS switched from commodity to R&E. If the populations experienced similar AS path lengths at the same configuration change points in each experiment, their lines would entirely overlap. In the SURF experiment (Figure 8a) we observed that the *Participant* (U.S. domestic R&E networks) class switched to R&E one prepend adjustment after the *Peer-NREN* class did. Those ASes likely learned routes with relatively long AS paths via Internet2, which required an additional commodity prepend before the R&E path was shorter. In the Internet2 experiment, the lines are similar, although twice as many *Peer-NREN* ASes switched at 2-0 than did *Participant* ASes. We believe the available commodity route had a longer AS path by the time it reached these international networks.

In Figure 7, the first row in case J demonstrates that networks which ignore AS path length and select the oldest route will switch from commodity to R&E with configuration 0-1. A total of 8 prefixes originated by 4 ASes switched at configuration 0-1 in both SURF and Internet2 experiments; all involved prefixes were Peer-NREN prefixes. We therefore believe that there is limited evidence that R&E ASes did not consider AS path length in route selection and broke ties with route age.

Ethics

We minimized the effect of our experiments by (1) re-using existing datasets that identify responsive systems in prefixes instead of conducting our own scans, (2) probing at a low rate (100pps) from the measurement host, with benign ICMP echo, TCP SYN, and UDP probes, (3) allowing at least an hour between changes of BGP configurations, (4) using a BGP prefix that does not host services affecting end users.