CICI:TCR: Routing Operations Observational Technology: Building to Ensure Efficacy of Research (ROOTBEER)

1 Introduction

Scientists, researchers, students, and educators access the Internet via *research and education (R&E) networks*, which include campus networks, regional networks, and national backbones. These R&E networks offer specialized capacity to their member institutions that is not available through commodity (commercial) networks, enabling data-intensive collaborative research and discovery across scientific disciplines. To support seamless workflows for domain scientists, R&E networks typically prioritize R&E routes ahead of commodity Internet routes. Although optimal for scientific collaboration, this approach introduces vulnerabilities to the integrity of the underlying routing infrastructure for two related reasons.

First, prioritizing R&E paths increases the importance of ensuring correctness in all R&E router configurations. Even minor misconfigurations have resulted in significant route leaks, inadvertently transmitting sensitive U.S. scientific traffic through unintended international routes [1]. Second, academic networks often operate under constrained budgets with limited operational staffing. Consequently, many of these networks have yet to implement routing security innovations recommended as best practices for over a decade [2,3]. Furthermore, as with many security measures, new routing innovations can introduce unforeseen vulnerabilities, creating additional barriers to their adoption.

To address these critical issues, we propose to develop a security-focused routing observatory and operational support system designed to ensure that routing policies align with the security and integrity goals of the U.S. science ecosystem. Recent developments in Internet address policy provide a timely opportunity for the academic community to take collective action to improve the resilience and trustworthiness of its networks. Our project is structured into three tasks, executed in close collaboration with Internet2 to ensure maximum effectiveness:

1. **Measurement and Analysis**: We will adapt and integrate recently developed measurement and analysis capabilities to detect route leaks between R&E and commodity (non-R&E) networks. Such leaks, whether due to accidental misconfigurations or malicious activities, threaten the availability and integrity of critical R&E network services.
2. **Operational Dashboard**: We will develop a user-friendly dashboard to operationalize measurement findings. This dashboard will leverage Internet2's established technical platforms, and enable a broader set of R&E community stakeholders to identify vulnerabilities, quantify network resilience gaps, and implement security best practices effectively, reducing the risk of misconfigurations.
3. **Community Engagement**: We will engage actively with the community to disseminate the innovations developed through this project, encourage widespread adoption, and quantitatively assess impact based on deployment metrics.

1.1 Intellectual Merit

Our project will enable the U.S. R&E networking ecosystem to leverage recent advances in network measurement and analytics, thereby providing a rigorous assessment of network infrastructure security properties otherwise unobservable by external parties. We incorporate FAIR principles (Findable, Accessible, Interoperable, and Reusable) into our cybersecurity innovations, facilitating creation of a robust routing security auditing framework that will safeguard the integrity of U.S. scientific workflows. Moreover, our measurement innovations are designed to improve *usability, reproducibility, and interoperability* of scientific software tools to study the Internet.

1.2 Broader Impacts

Integrating these security-focused innovations into the academic network operations community will significantly enhance the community's capacity to adopt and maintain critical cybersecurity best practices, addressing historical barriers to widespread implementation. Aligned with the objectives of the Cybersecurity Innovation for Cyberinfrastructure (CICI) program, our project will empower network operators to proactively secure the integrity, availability, and performance of the U.S. scientific cyberinfrastructure. By fostering broader community engagement, increasing operational transparency through intuitive dashboards, and enabling rigorous monitoring capabilities, our approach will ensure sustained improvements in network resilience, reliability, and trustworthiness. Ultimately, these advancements will support uninterrupted scientific collaboration and innovation across research domains.

2 Why this project aligns with CICI:TCR solicitation criteria

Our project targets the objectives described in NSF's Cybersecurity Innovations for Cyberinfrastructure solicitation, and in particular its *Transition to Cyberinfrastructure Resilience* track.
Objective: This project promises to improve the robustness, trustworthiness, integrity, and/or resilience of the networks underpinning all U.S. scientific cyberinfrastructure.
Approach: Our approach integrates testing, evaluation, hardening, validation, and technology transition of recently designed and prototyped active measurement capabilities (§3). This project prioritizes transition of these capabilities to improving the security posture of U.S. scientific research infrastructure, namely Internet2 and the campus and research networks it supports.

Specifically, we will create a framework for independent validation of routing security configurations to support improved CI security. The outcomes will achieve not only the primary benefit of directly improving the security of scientific CI, but also the secondary benefit of demonstrating how cybersecurity innovations can transition into operational practice. Moreover, we structure the project to include a pathway for sustainability by Internet2 beyond the life of the grant. Our proposed transition of monitoring capabilities to the R&E community will extract actionable insights into routing security vulnerabilities of scientific infrastructure, to safeguard scientific workflows for domain scientists.

2.1 Scientific infrastructure and environments that will benefit

This project targets the U.S. R&E network community and the scientific users it serves, which includes campus networks, regional networks, and the U.S. national R&E backbone (Internet2). The platform resulting from this project will serve as a role model for other regions, so although our primary target is the U.S. scientific CI community, the global CI community will also benefit.

2.2 How this project will benefit scientific communities

This project will advance the ability of R&E networks to detect routing misconfigurations that expose scientific workflows to tampering, redirection, or other disruption. Moreover, our measurement innovations are designed to improve *usability, reproducibility, and interoperability* of software to support active Internet measurement experiments. By facilitating unambiguous labeling of measurement data, the tools will support a scientific approach to infrastructure security evaluations, thus facilitating objective measurement of the uptake of improved security practices.

2.3 Threat model motivating proposed solutions

Our threat model is accidental or malicious configuration of routing that has the ability to induce disruption or misdirection of data flows. Steven Wallace, Director of Routing Integrity for Internet2 and leading the Internet2 team in this collaboration, recently described three impactful routing security incidents in 2024 [1]. The three incidents exemplified the three most common threats. In the first, commercial routes leaked into the R&E ecosystem, which induced delays and outages for scientific CI. In the second incident, R&E routes leaked to commercial providers, which caused traffic from major cloud providers to route return traffic through the Middle East to reach a U.S. university. The third incident was a route hijack targeting a U.S. R&E regional network, which rendered key services unavailable for networks who had not implemented RPKI-ROA technology [4].

2.4 Unique properties of R&E network infrastructure that motivate our work

The research and education (R&E) networking ecosystem exhibits distinctive routing policies that introduce unique vulnerabilities, motivating our proposal. At the core of this ecosystem are backbone networks such as Internet2, which serve as critical interconnection fabrics for regional and national R&E networks worldwide. These backbone networks treat connected institutions as routing customers, propagating their routes to other R&E peers like GÉANT and AARNet, thus facilitating global connectivity among research and educational entities.

Members directly connected to these backbone networks often have their own downstream customers and peers, typically preferring routes originating from their immediate networks over those provided by backbone providers. However, unlike conventional commercial routing, the R&E ecosystem permits and actively facilitates the exchange of peer routes among backbone networks. For instance, Internet2 routinely exports routes from one national R&E network (NREN) peer to another, effectively establishing a unified global research network.

Furthermore, participants within the R&E community consistently prioritize routes within this ecosystem over commodity Internet paths, ensuring optimal use of dedicated R&E capacity. While advantageous for efficiency and collaboration, this uniquely permissive and interconnected routing environment significantly amplifies the risk posed by routing misconfigurations. Errors such as accidental leaks of commercial routes into the R&E space can rapidly propagate globally, unintentionally redirecting critical scientific traffic through unintended and potentially insecure paths. Indeed, past incidents have resulted in sensitive U.S. scientific data inadvertently traversing networks in distant regions, including the Middle East [1]. This amplified vulnerability underscores the necessity of our proposed effort to rigorously understand and mitigate these unique routing risks within the global R&E networking infrastructure.

2.5 Software licensing plans

The project builds on existing software base that uses a GPL license, and any modifications to that software will respect the GPL. Other software will use open source licenses as approved by UC San Diego and/or Internet2 as appropriate.

2.6 Transition and sustainability plan

Internet2 intends to integrate the newly developed monitoring capabilities seamlessly into its existing Routing Integrity tools within the Internet2 Insight Console. The transition will leverage Internet2's established technical platforms and operational frameworks, ensuring sustainability and scalability. Ultimately, these new tools will become integral components of Internet2's sustained commitment to providing actionable intelligence aimed at enhancing global routing security.

2.7 Ethical and operational concerns

Although some in the R&E community are aware of the unique routing policies that characterize networks, there is scant appreciation for the operational risks inherent to the model. One outcome of this project is to raise awareness of these risks and offer infrastructure support to mitigate them.

Although our measurements do not require explicit consent of networks we measure, we take seriously the potential exposure of measured networks to compromise and undertake responsible disclosure of potential vulnerabilities directly to those networks.

2.8 Partnerships and collaborations

This project will strengthen and extend a collaboration with Internet2, and expand it to benefit other U.S. R&E regional networks, Internet2 will engage closely with identified transition partners and stakeholders throughout the deployment process to ensure alignment with community needs and maximize impact. Our recent presentation to the Internet2 Tech Ex community [5] reinforced our position as a trusted source of measurement data and analysis for the R&E community. We hope to present at future Tech Ex and Campus Cyberinfrastructure PI meetings (§4.3.1).

2.9 Quantitative evaluation metrics

Quantitative metrics for transition success will include an increased number of R&E address blocks registered in authenticated IRRs, and measurable improvements in routing security indicators, (e.g., reduction in routing anomalies, enhanced RPKI validation coverage), and expanded usage and utility of routing data analytics across the community. RIPE's data archive [6] will facilitate evaluation of whether our project correlates with increased routing security adoption rates for R&E networks over time.

3 Description of Innovation to Transition

In this section we provide details on a specific use case of our recent cybersecurity innovations to measure and analyze the risks outlined in the previous section. This measurement experiment represented a successful example of modern, complex, data-driven, distributed, rapid, and collaborative science on the backbone network supporting the U.S. scientific research community.

3.1 Method: Inferring R&E route preference

In collaboration with Internet2, we developed and validated a novel active measurement method that inferred route preference of R&E networks, and investigated the degree to which R&E members preferred available R&E routes to ensure R&E members use available R&E capacity. Recall from §2.4 that this feature of R&E networks amplifies the impact of unintentional leaks.

It is difficult for a third-party operator or researcher to determine relative route preference of R&E institutions such as universities, as existing data-driven methods for ascertaining or inferring route preference have limited coverage. The first method – using public BGP data to observe if members select R&E routes – requires a view from each member. As of November 2024, fewer than 1% of all ASes (≈400 ASes) provided a full BGP routing table to RouteViews or RIPE RIS. Of the 2,659 R&E ASes we examined in November 2024, 27 (1.0%) provided a public BGP view. The second method – using looking glasses to observe localpref values assigned to routes – also had few ASes publicly providing that capability. In 2023, Kastanakis *et al.* found 10 ASes with looking glasses from an initial list of 76 that provided localpref values [7]. The third method – using public

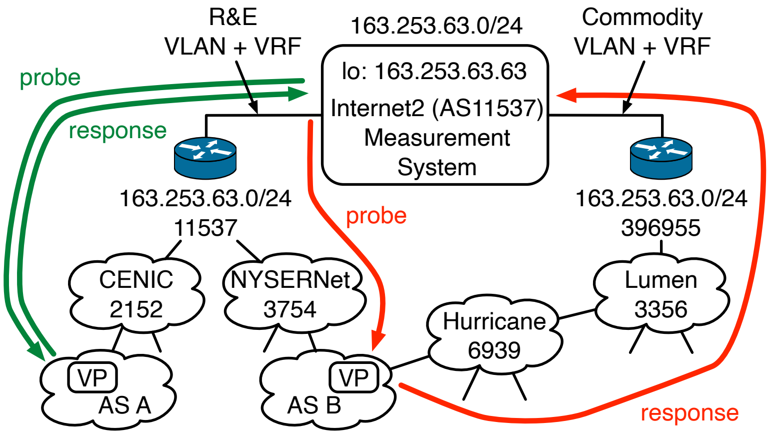

Figure 1: Internet2 measurement system configuration. The system was connected to two VRFs that had different export policies to R&E and commodity networks. We inferred the route used by member networks based on the VLAN interface on which the response arrived.

vantage points (VPs) to collect traceroute paths – also requires wide deployment of VPs. As of November 2024, RIPE Atlas had VPs in 3,808 networks (5.0% of all routed ASes) – with VPs in 215 (8.1%) of 2,659 R&E ASes. The fourth method – using Routing Policy Specification Language (RPSL) databases where operators *may* encode localpref in rules – suffers from inconsistent coverage across RIR regions and staleness of data [8]. As of November 2024, the RIPE database had 39,337 ASes, 7,331 of which included at least one rule with localpref. The RIPE database represented nearly half of the ASes and 73.3% of data by size in RPSL across 13 IRR databases considered by He *et al.* in June 2023 [8]. In 2023, Kastanakis *et al.* reported variance between deployed localpref and "abstract policies described for documentation purposes" in RPSL [7].

The new method we developed involved probing addresses in R&E prefixes from a single system in Internet2, configured as illustrated in Figure 1. We assigned a publicly-routed IPv4 address within a measurement prefix to the loopback interface on Internet2's measurement system, and used that source address in our probe packets. The measurement system was connected to one router via two VLAN interfaces. Internet2 has separate virtual routing and forwarding (VRF) instances for commodity and R&E routing, which the router presented to the measurement system as separate VLAN interfaces. The VLAN on which a response packet arrived identified the VRF through which the packet flowed, and therefore whether the response followed R&E or commodity routing. Consequently, *every responder became a VP that we could use to infer the routing policy of the network hosting the responder*. We developed custom software to run on Internet2's measurement system, using CAIDA's python module [9] to coordinate measurements with scamper [10]. We extended scamper to use the IP_PKTINFO ancillary message [11] via the *recvmsg* system call to obtain the interface on which the operating system received response packets. If responses arrived on the measurement system's R&E interface, then we inferred that the member had selected an R&E route, otherwise they had selected a commodity route.

3.2 Design of experiment to ascertain R&E route preference

To support our experiment, we BGP-advertised our measurement prefix with origin AS 11537 to R&E networks, and with origin 396955 to commodity networks via Lumen (AS 3356); these announcements were covered by RPKI ROAs. We conducted a series of nine experiments, using a different AS prepend configuration in each, to infer member sensitivity to AS path length. We use "4-0" to refer to the experiment with 4 prepends of the R&E ASN and no prepending of the

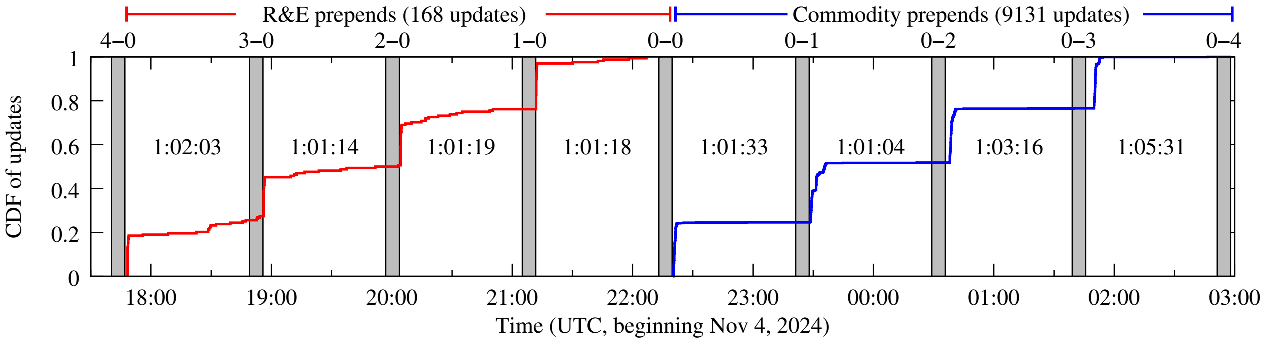

Figure 2: Experiment flow between active probing windows (≈7 minutes, grey bars) and measurement prefix BGP activity during our experiments. In the first five (R&E prepends) experiments, updates after R&E convergence exclusively occurred on commodity routes, and were 40 of the 168 updates during the R&E prepends phase.

Inference	Prefixes		ASes	
Always R&E	10,027	81.3%	1,925	74.1%
Always commodity	756	6.1%	356	13.7%
Switch to R&E	1,147	9.3%	336	12.9%
Switch to commodity	2	0.0%	1	0.0%
Mixed R&E + comm.	393	3.2%	254	9.8%
Oscillating	6	0.0%	4	0.2%
Total:	12,331		2,597	

Table 1: Results for tested prefixes. VPs in most (10,783, 87.4%) prefixes were insensitive to our AS path length changes, and most (10,027, 81.3%) always used the R&E route.

commodity ASN, “0-0” to refer to no prepending of either R&E or commodity ASN, and “0-4” to refer to the experiment with 4 prepends of the commodity ASN and no prepending of the R&E ASN. The order of our nine experiments 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. 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 suppression of our BGP configurations [12, 13]. To minimize the effect of RFD but still allow our experiments to complete within a work day, we conducted active probing one hour after changing the configuration of our BGP announcements. In 2020, Gray reported that ≈9% of the ASes they measured enabled RFD, few ASes damped prefixes longer than 15 minutes, and they did not observe suppress times longer than one hour [14].

3.3 Experiment: execution

On November 4th 2024, we surveyed 18,953 prefixes originated by 2,659 R&E member ASes, obtaining responses from VPs within 12,331 (65.1%) prefixes across 2,597 (97.7%) member ASes. Table 1 summarizes our experiment results. We characterized prefixes that had a response from at least one VP during every active probing round, so this table excludes 142 of 12,473 prefixes that we probed. Overall, VPs in most prefixes were insensitive to our BGP configuration changes – response traffic for 81.3% of prefixes always arrived via the R&E route, and always via commodity for 6.1% of prefixes. VPs in 9.3% of the prefixes switched to the R&E route when we used a prepend configuration that would cause a router on the return path to prefer the R&E route – implying that that router's BGP instance used AS path length (and not localpref) to select the route. A single AS switched from R&E to commodity during our experiment, even as we increased commodity prepends. We discussed our findings with an operator at the AS, who reported that an outage during our experiment caused their route to our measurement system to be over commodity.

3.4 Experiment Validation against BGP data

Of the 2,597 ASes with at least one prefix responsive to our method, 27 ASes also provided a public BGP view to RouteViews and/or RIPE RIS collectors, which we compared to our prefix-level inferences. For each collector, we downloaded the November 4th 16:00 UTC RIB file and all update files through the entirety of our experiment, extracting activity for our measurement prefix (Figure 1). Most ASes had active measurement inferences for multiple prefixes; if an AS had different inferences for different prefixes, we considered the most frequent inference. Two ASes had no most-frequent inference and we did not include them in this validation. For the remaining 25 ASes, we say that the active measurement inference is consistent with the public BGP view if the origin AS or ASes observed in BGP are expected given the inference. For example, we would expect to only observe routes for our measurement prefix originated by AS11537 in BGP if we inferred through active measurement that responses to our probe packets used the R&E route.

Overall, 23 of 25 ASes had BGP activity congruent with active measurement inferences. For two other 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 two incongruent ASes and received responses from both. Both reported that their 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, because public BGP collectors request complete routing tables from their peers. That is, our policy inference was correct.

3.5 Experiment validation against operator ground truth

We contacted operators at nine ASes with our findings, and received responses from six, covering the spectrum of our inferences. For two ASes, we had inferred that they assigned equal localpref to both R&E and commodity routes, as we had inferred their return route was sensitive to AS path length, which they confirmed. For another AS, we had observed one /24 prefix where two systems replied via R&E, while a third system's reply had returned via commodity. This AS operator confirmed that systems in prefixes originated by their AS would use the R&E route, but that we had probed a router that used an address from their prefix for interconnection, and that this router did not have an R&E route. Finally, we heard from the other three ASes that our inferences that they preferred R&E routes over commodity routes were correct.

Overall, our preliminary work conclusively demonstrated two important conditions. First, *responsive systems in R&E networks can be used as VPs to infer routing policy.* Second, *R&E members almost always preferred available R&E routes*, creating the unintended risk of dangerous leaks. That is, because global R&E network participants (2,659 ASes in November 2024) prefer R&E routes, any leak of a commodity route will become the preferred route for nearly the entire R&E ecosystem distributed across the globe. Accidental misconfigurations have created route leaks that led to transmission of U.S. scientific traffic through unintended international routes [1].

4 Infrastructure development agenda: preventing diversion of scientific data

The new measurement technique and its application described in §3 demonstrated not only that R&E network operators have a legitimate concern about potential route leaks, but also that we have a rigorous and scalable way to detect them. We now propose a monitoring capability that will *rapidly detect diversion of scientific data* to networks the data should not cross. We structure our project in three tasks. Our first task will adapt and integrate recently developed monitoring capabilities to detect routes leaked between R&E and commodity networks. These route leaks are caused by misconfigurations and malicious attacks that can pose to R&E network connectivity and services [1]. Our second task is to create a dashboard to visualize the results of our measurements to support an expanded set of R&E community stakeholders to improve their routing security posture. This dashboard will quantify and visualize gaps in R&E infrastructure resilience, and facilitate deployment of security best practices while minimizing the risk of misconfigurations. Our third task is community engagement to socialize the innovations we have integrated, and evaluate their impact.

4.1 Task 1: Monitoring to detect route leaks to/from R&E networks

Rather than probing only R&E routes as in §3, the monitoring infrastructure we create for this project will actively probe IP addresses from three distinct categories:

1. IP addresses within the known Research and Education (RE) infrastructure
2. IP addresses outside the R&E infrastructure (e.g., commercial or commodity networks)
3. IP addresses announced to Internet2 by National Research and Education Network (NREN) peers, but not yet classified as known R&E infrastructure and currently rejected by Internet2

Probing IP addresses within the known R&E infrastructure will verify whether networks are accurately configured to route traffic exclusively over trusted R&E links or if some traffic might inadvertently traverse commercial networks. Additionally, the geographically distributed nature of the probe receiver agents will help identify whether traffic paths adhere to specific regulatory or data sovereignty requirements by remaining confined to appropriate segments of the R&E infrastructure.

Probing IP addresses outside the recognized R&E infrastructure will detect instances where R&E network routes have unintentionally leaked into commercial networks, resulting in unauthorized commercial traffic crossing R&E infrastructure. Distributed probe receiver agents will further facilitate pinpointing the origins of these route leaks.

Finally, probing IP addresses that are announced to Internet2 from NREN peers but not yet classified as known R&E infrastructure will support the evaluation process to determine whether these routes represent legitimate expansions of the RE infrastructure or if they are unintended leaks requiring remediation.

4.1.1 Detecting R&E routes leaked into commodity networks

First, we will identify inadvertent route leaks of R&E address space into commodity networks. The experiment shares concepts with our previous work described in §3 for inferring R&E route preference [5], where we announced routes over both R&E and commodity networks, and inferred properties of the network hosting the VP based on the interface that receives a response. Figure 3a illustrates the design.

In these experiments, rather than limiting the IP addresses probed (the VPs) to being within R&E prefixes, we will probe IP addresses that reside in all Internet-routed prefixes, as we seek to identify commodity networks that have received R&E routes that they should not have. We will

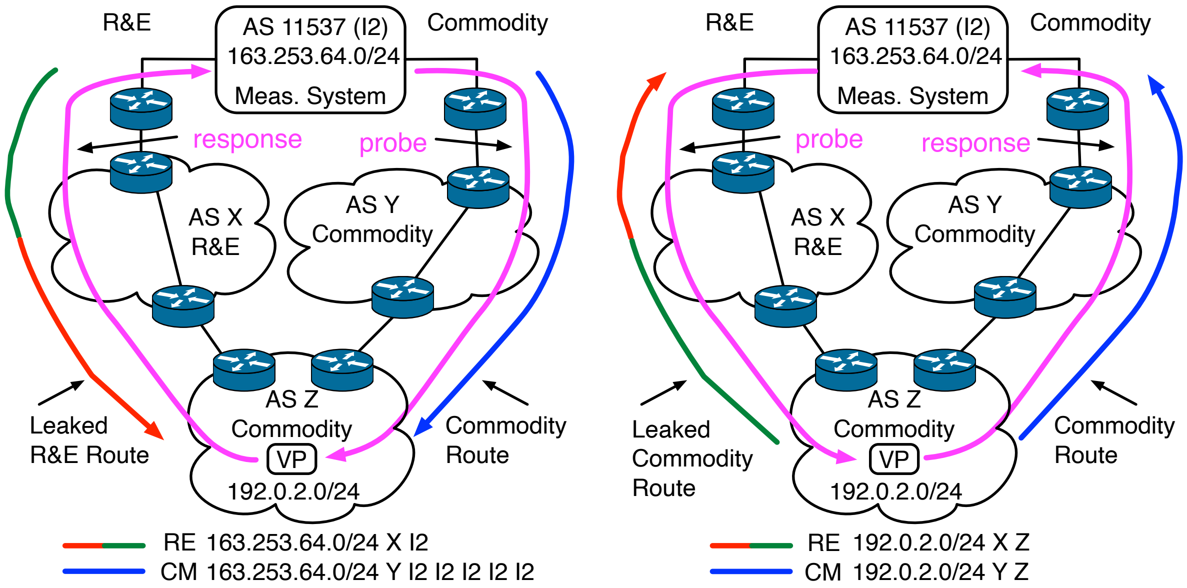

(a) Leaks from R&E into commodity

(b) Leaks from commodity into R&E

Figure 3: Inferring leaks between R&E and commodity ASes. In (a), a probe to commodity network Z will return to Internet2 over R&E if X leaked an R&E route to Z, because Z's route via X is the shortest Z's available AS paths for the prefix. This pattern indicates that X leaked an R&E route to commodity. In (b), a probe to commodity network Z sent via X (R&E) will return to Internet2 over commodity, because Z does not have an available R&E route. This indicates that X leaked a commodity route to R&E.

use ISI's Internet Addresses IPv4 Response History Dataset [15], which summarizes previous ISI censuses of the IPv4 address space [16], to identify responsive addresses to probe. In addition, rather than *vary* AS path prepending to identify R&E networks that assign equal preference to routes received from both R&E and commodity providers, we will announce an R&E route with no prepending, and a heavily prepended commodity route. A commodity network that receives the route leaked by an R&E member will prefer the R&E route because it has the shorter AS path, so a probe to a commodity prefix will return over the R&E interface.

There are two possible explanations for this behavior. The first is that the R&E table lacks R&E prefixes, perhaps due to routes from R&E members that Internet2 currently filters because they are not documented by partner networks as being from an R&E institution. We can use manual inspection to account for these prefixes. The more likely explanation is that a commodity network received an R&E route via a leak. Figure 3a implies that we can identify the AS (X) that leaked the route to the commodity network. Unfortunately, it is not that simple, as we can only identify that the route was returned over an R&E path. For example, the leaking AS may be a European AS many ASes downstream from GÉANT.

We will therefore extend our existing Internet2-hosted active measurement platform (§3) to support BGP-anycasting measurement prefixes. This capability will enable us to narrow down the possible ASes from where the BGP leak of R&E space could be occurring. Internet2 will work with R&E partners in the U.S. (state networks) and globally (other NRENs) to BGP-announce the measurement prefix, so that R&E institution closest (in BGP hops) to the leak will receive responses to our active measurements. The partners will work with Internet2 to establish software tunnels (GRE, or other technology) to relay packets to the Internet2-hosted active measurement platform hosting our software. The interface corresponding to the partner software tunnel on which Internet2 receives the packet identifies the partner institution closest to the leaking party.

4.1.2 Detecting commodity routes leaked into R&E networks

Because the R&E ecosystem prefers routes in the R&E ecosystem over other commodity routes, inadvertent leaks of commodity routes into R&E can result in significant volumes of traffic forwarded to R&E institutions without the capacity forward the traffic, harming their ability to forward R&E traffic. Internet2 has recently announced a policy where they will soon only permit prefixes within a defined list from being routed by their network [17]. However, this policy can inadvertently cause Internet2 to reject R&E prefixes, if those members did not proactively advise their upstream networks of their intention to use new address space. Filtering R&E prefixes reduces availability of paths to these prefixes, and is contrary to Internet2's goal of carrying R&E traffic over R&E infrastructure. With the proposed innovations integrated into Internet2's monitoring infrastructure, it will be safe for Internet2 to launch this new routing policy.

We will integrate available indicators to assist operators in triaging these prefixes, identifying prefixes that are most likely R&E prefixes (Internet2 operators should confirm these with members) and prefixes that likely leaked (Internet2 operators should advise their members of the misconfiguration). First, we will use daily MAnycastR anycast census [18] to identify members leaking prefixes that are anycast. Anycast prefixes are designed to localize traffic, and their appearance in R&E routing means that R&E institutions around the world could erroneously use that anycast instance instead of instances local to those institutions. The University of Twente, supported by CAIDA's Archipelago measurement infrastructure, has operationalized a daily MAnycastR census, which is updated daily.

Second, we will use our Internet2-hosted active probing infrastructure to identify prefixes that are likely commodity, and should be filtered. Figure 3b illustrates the idea: if Internet2 receives

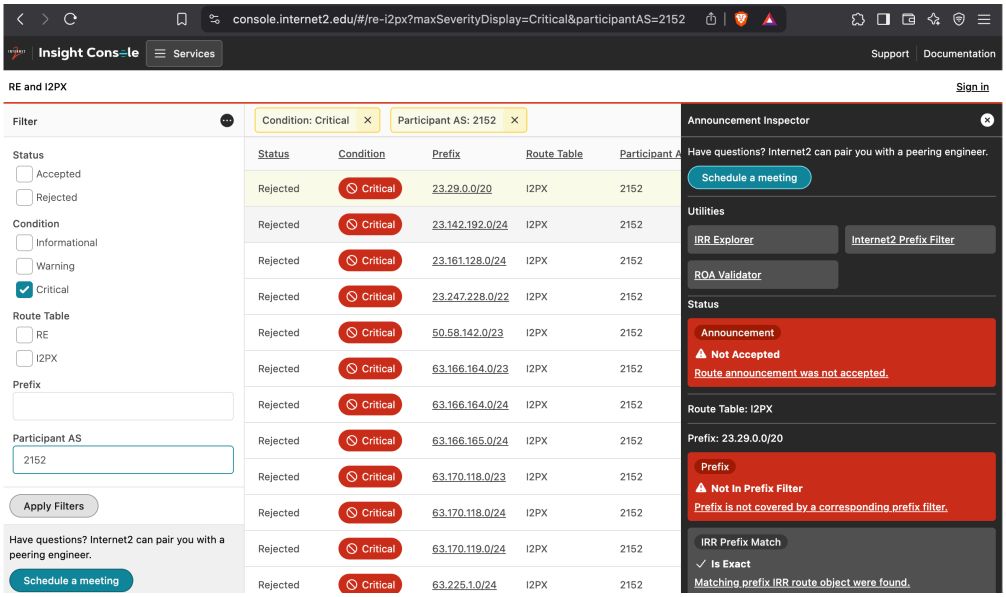

console.internet2.edu/#/re-i2px?maxSeverityDisplay=Critical&participantAS=2152

RE and I2PX

Filter

Status
□ Accepted
□ Rejected

Condition
□ Informational
□ Warning
☑ Critical

Route Table
□ RE
□ I2PX

Prefix

Participant AS
2152

Have questions? Internet2 can pair you with a peering engineer.

Condition: Critical ×
Participant AS: 2152 ×

Status	Condition	Prefix	Route Table	Participant A
Rejected	Critical	23.29.0.0/20	I2PX	2152
Rejected	Critical	23.142.192.0/24	I2PX	2152
Rejected	Critical	23.161.128.0/24	I2PX	2152
Rejected	Critical	23.247.228.0/22	I2PX	2152
Rejected	Critical	50.58.142.0/23	I2PX	2152
Rejected	Critical	63.166.164.0/23	I2PX	2152
Rejected	Critical	63.166.164.0/24	I2PX	2152
Rejected	Critical	63.166.165.0/24	I2PX	2152
Rejected	Critical	63.170.118.0/23	I2PX	2152
Rejected	Critical	63.170.118.0/24	I2PX	2152
Rejected	Critical	63.170.119.0/24	I2PX	2152
Rejected	Critical	63.225.1.0/24	I2PX	2152

Announcement Inspector

Have questions? Internet2 can pair you with a peering engineer.

Utilities

Status

Announcement
Not Accepted
Route announcement was not accepted.

Route Table: I2PX

Prefix: 23.29.0.0/20

Prefix
Not In Prefix Filter
Prefix is not covered by a corresponding prefix filter.

IRR Prefix Match
✓ Is Exact
Matching prefix IRR route object were found.

Figure 4: Snapshot of Internet2 routing knowledge dashboard into which we will integrate new monitoring innovations

a route for a prefix not in its allowed list, and systems in that prefix choose a commodity return route, then the announced prefix is unlikely to be R&E.

4.2 Task 2: Making results accessible and actionable

Our second task is to integrate the results of our new measurement capability into a dashboard to serve the global R&E community. As part of its Routing Integrity initiative, Internet2 currently maintains and operates several critical infrastructure systems that collect and integrate heterogeneous datasets, including BGP data (from Internet2 itself as well as RouteViews and RIPE), WHOIS address registration data from ARIN, Internet Routing Registry (IRR) data, and Resource Public Key Infrastructure (RPKI) data from multiple trusted sources.

We will synthesize these data into a comprehensive routing knowledge board, facilitating improved routing security posture and informed decision-making within the research and education community. Internet2 will integrate the newly developed capabilities seamlessly into its existing Routing Integrity tools within the Internet2 Insight Console. The transition will leverage Internet2's established technical platforms and operational frameworks, ensuring sustainability and scalability.

4.2.1 Making measurement results accessible to community via knowledge dashboard

Internet2's current routing knowledge dashboard (Figure 4) displays all routes announced to Internet2 by its regional networks prior to the application of routing policies. Consequently, some routes shown may ultimately be rejected by policy. For example, a route inadvertently leaked from an Internet transit provider would appear with a status indicating rejection. The "Announcement Inspector" provides detailed, actionable insights for individual network prefixes. Network engineers can leverage this tool to assess routing security attributes, such as the presence of RPKI Route Origin Authorizations (ROAs) and matching Internet Routing Registry (IRR) route objects. While this information clarifies how the Internet2 network forwards traffic toward specific destinations,

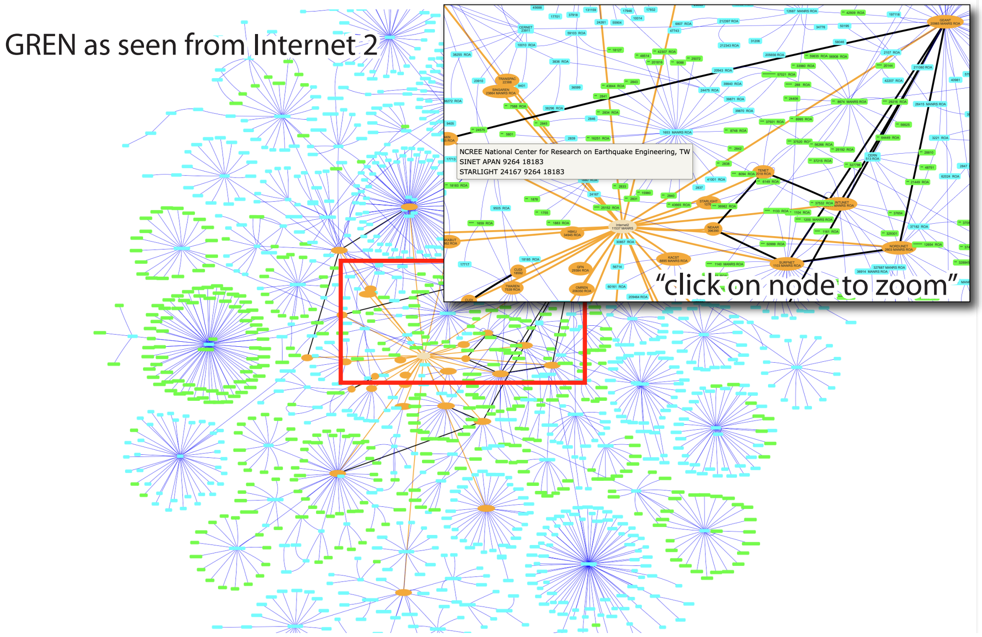
GREN as seen from Internet 2
NCREE National Center for Research on Earthquake Engineering, TW
SINET APAN 9264 18183
STARLIGHT 24167 9264 18183
Internet2 11537 MANRS
“click on node to zoom”

Figure 5: Sample Snapshot of Global R&E network topology. Dots in green have multiple routes. Overlay is a zoomed-in view of connectivity from research center in Taiwan.

merely knowing the announced routes does not indicate whether the originating network will route its outbound traffic via Internet2 or through alternative transit providers.

Through the proposed enhancements, the Announcement Inspector will incorporate additional data reflecting each regional network's effective outbound routing policies. Specifically, it will verify if networks are appropriately configured to direct their traffic over the R&E infrastructure, thereby providing a more comprehensive view of actual network routing behaviors.

In addition to augmenting the existing dashboard with regional networks' outbound routing policies, Internet2 will introduce an entirely new dashboard dimension that reflects similar information for the Global R&E infrastructure, facilitated by the various National Research and Education Networks (NRENs) peering with Internet2. This will enable network engineers in the U.S. to determine the likely path that traffic from remote sites will take. With the addition of IP anycast receivers for the active measurements (§4.1), engineers will not only confirm if the path traverses the R&E infrastructure versus commercial transit but also identify which intermediate R&E transit networks are likely to carry the traffic. This new dimension will also highlight which routes from NRENs are rejected because they originate from commercial networks that are not intended to be NREN-accessible.

4.2.2 Visualization of AS topology

CAIDA and Internet2 will develop an NREN Peering Dashboard widget to visualize the entire BGP topology encompassing the 2,700 Autonomous System Numbers (ASNs) within the global R&E network, as viewed through Internet2's peering relationships. Figure 5 shows an example of the global R&E BGP topology presented in a zoomed-out view and then zoomed in to show detail on multiple routes to a research center in Taiwan (based on a working prototype). The active measurement data we collect in could allow us to augment, refine, and validate this BGP topology view.

4.2.3 RPKI-ROA deployment support

Many national and regional R&E networks hold IP address allocations that they reallocate or reassign to their members. For example, Internet2 obtained an IPv6 /32 allocation in 2001, portions of which they reallocated to over 20 of its members over the next decade. More common are IPv4 /16s held by regional networks that they reassign to their members. While the original holders of these IP address allocations are the only entities authorized to generate Resource Public Key Infrastructure Route Origin Authorizations (RPKI-ROAs) to secure these addresses, they often lack complete or current documentation detailing how these sub-allocations are used and announced on the Internet. This information gap hinders the creation of accurate and effective RPKI-ROAs.

Internet proposes to build an RPKI-ROA planner to address this challenge by aggregating critical information from multiple authoritative sources: ARIN's WHOIS database, historical BGP data from RIPE, Internet Routing Registry (IRR) records specifying authorized Autonomous System Numbers (ASNs), and and NSRC's current data on announcements among other R&E networks. Since the ROA creation sequence for overlapping IP address spaces can significantly impact network stability, the Planner will also recommend an optimal order for creating these authorizations. By consolidating this information, the Planner will equip IP address holders with comprehensive knowledge required to create precise RPKI-ROAs for their reallocations and reassignments.

4.3 Task 3: Community engagement and evaluation of impact

Our third task is community engagement and evaluation of impact. The goal of this task is to socialize and expand use of the dashboard, evaluate the impact of NSF's investment into this project using quantitative metrics, and demonstrate the use of this platform to monitor the risks of impending Internet governance changes that hold particular resilience risk for the R&E cyberinfrastructure ecosystem.

4.3.1 Outreach to explain and expand use of dashboard

Community engagement will include quarterly zoom calls, annual workshop events, and presentations (if/as invited) at campus cyberinfrastructure events to socialize the innovations we have integrated, and report quantitative metrics of impact of its deployment. This project is inspired in part by a recent successful collaboration with Internet2 network engineers, and is structured to deepen this collaboration with Internet2 and expand it to benefit other U.S. R&E regional networks, starting with collaborators at CENIC, CAAREN, ONENet, Link Oregon (see LOCs).

4.3.2 Evaluation of impact

As mentioned in §2.9, we will rely on quantitative metrics to gauge the success of this transition project. We will measure the increase in R&E address blocks registered in authenticated IRRs and RPKI, and other measurable improvements in routing security indicators, e.g., reduction in routing anomalies, enhanced RPKI validation coverage. RIPE's data archive [6] will facilitate evaluation of whether our project correlates with increased routing security adoption rates for R&E networks.

4.3.3 Broader impacts: monitoring risks to security and resilience in R&E networks

A critical advancement in routing security policy involves requiring Internet Exchange Point (IXP) members and their downstream customers to ensure prefixes announced to IXP route servers are validated through either Resource Public Key Infrastructure (RPKI) Route Origin Authorizations (ROAs) or records in an authenticated Internet Routing Registry (IRR). These authenticated databases are typically maintained by Regional Internet Registries (RIRs) such as ARIN, APNIC, and RIPE NCC.

In North America, to leverage these security mechanisms, institutions must hold an active service agreement with ARIN. Historically, universities and the U.S. federal government, holding legacy IP address resources, have hesitated to enter into these agreements. However, significant changes in ARIN's fee structure at the end of 2023 incentivized most U.S. R&E institutions to finally secure agreements enabling authenticated registrations.

Leaders at Internet2 successfully encouraged ARIN to publicly disclose service agreement statuses, facilitating targeted outreach efforts. As of March 2025, 93% of U.S. R&E institutions now possess the capability for secure registration, yet adoption remains inadequate. Of the approximately 6,500 prefixes originating within the U.S. R&E community, nearly 6,000 are eligible for secure registration; however, fewer than 1,400 currently have validated records.

As IXPs increasingly adopt this new filtering standard—requiring authenticated records for route propagation—R&E and campus networks lacking proper registrations face tangible risks. These networks will experience reduced resilience and degraded performance due to limited routing paths, significantly impacting their ability to support critical scientific research activities.

The comprehensive monitoring infrastructure we develop will allow for a precise evaluation of the operational impacts of this policy shift on the R&E ecosystem. Leveraging data from PeeringDB, we will systematically identify relevant IXP route servers. Complementary routing data from Packet Clearing House (PCH) will allow us to track and analyze the prefixes propagated through these IXPs. Preliminary analyses reveal notable exposure: Equinix route servers at strategic locations – Ashburn, Chicago, and Dallas – carry the highest concentration of R&E prefixes, each hosting at least 1,500. However, only approximately 300 of these prefixes are currently validated by authenticated ARIN IRR records. Notably, Hurricane Electric announces nearly all these R&E prefixes at these major exchanges and nearly 200 additional exchanges worldwide, emphasizing the broad and urgent need for action.

We include this example of a broader impact of our proposed monitoring framework, which will deliver essential insights beyond the life of this project, empowering the R&E community to mitigate risks proactively and preserve the integrity and reliability of critical network-dependent research activities.

References

[1] Steve Wallace, “What the Research and Education Community Learned From Three Impactful Routing Security Incidents in 2024,” Sept. 2024. `https://internet2.edu/what-the-research-education-community-learned-from-three-impactful-routing-security-incidents-in-2024/`.

[2] D. Clark, C. Testart, M. Luckie, and kc claffy, “A path forward: Improving Internet routing security by enabling zones of trust,” *Journal of Cybersecurity*, 2024.

[3] B. Du, C. Testart, R. Fontugne, G. Akiwate, A. C. Snoeren, and k. claffy, “Mind Your MANRS: Measuring the MANRS Ecosystem,” in *ACM Internet Measurement Conference*, 2022.

[4] Internet2, “Protect Your Network With RPKI,” 2024. https://internet2.edu/network/initiatives-partnerships/routing-integrity/protect-your-network-with-rpki/.

[5] M. Luckie, K. Newell, S. Wallace, J. Bartig, J. Deaton, and kc claffy, “Inferring Relative Route Preference using Active Probing,” in *Internet2 Technical Exchange Meeting*, Dec. 2024. `https://users.caida.org/~mjl/tmp/lpp.pdf`.

[6] RIPE NCC, “RIPEstat Docs: RPKI History,” 224. `https://stat.ripe.net/docs/02.data-api/rpki-history.html`.

[7] S. Kastanakis, V. Giotsas, I. Livadariu, and N. Suri, “20 years of inferring inter-domain routing policies,” in *IMC*, Oct. 2023.

[8] S. He, Ítalo Cunha, and E. Katz-Bassett, “RPSLyzer: Characterization and verification of policies in internet routing registries,” in *IMC*, pp. 365–374, Oct. 2024.

[9] M. Luckie, S. Hariprasad, R. Sommese, B. Jones, K. Keys, R. Mok, and K. Claffy, “An integrated active measurement programming environment,” in *PAM*, Mar. 2025.

[10] M. Luckie, “Scamper: a scalable and extensible packet prober for active measurement of the Internet,” in *IMC*, pp. 239–245, Nov. 2010.

[11] W. Stevens, M. Thomas, E. Nordmark, and T. Jinmei, “Advanced sockets application program interface (API) for IPv6,” May 2003. `https://www.rfc-editor.org/rfc/rfc3542`.

[12] C. Villamizar, R. Chandra, and R. Govindan, “Bgp route flap damping,” Nov. 1998. `https://www.rfc-editor.org/rfc/rfc2439`.

[13] R. Bush, C. Pelsser, M. Kuhne, O. Maennel, P. Mohapatra, K. Patel, and R. Evans, “RIPE routing working group recommendations on route flap damping,” Jan. 2013. `https://www.ripe.net/publications/docs/ripe-580/`.

[14] C. Gray, C. Mosig, R. Bush, C. Pelsser, M. Roughan, T. C. Schmidt, and M. Wahlisch, “BGP beacons, network tomography, and bayesian computation to locate route flap damping,” in *IMC*, pp. 492–505, Oct. 2020.

[15] USC/LANDER project, “Internet addresses ipv4 response history dataset, predict id: Usc-lander/internet_address_history_it108w-20240711/rev14937,” July 2024.

[16] J. Heidemann, Y. Pradkin, R. Govindan, C. Papadopoulos, G. Bartlett, and J. Bannister, “Census and survey of the visible Internet,” in *IMC*, pp. 169–182, Oct. 2008.

[17] Steve Wallace and Thomas Araneta and Jeff Bartig, “Internet2 Routing Policy for International Research and Education Peers,” Dec. 2024. `https://docs.google.com/document/d/12z-019wD73aKcrhj0zCDqCOGvCN42tRLgdhKWhWS3C0`.

[18] University of Twente, “Manycastr anycast census,” Mar. 2025. `https://github.com/anycast-census/anycast-census`.