DND-Db: A Democratized Network Data Database for Tailored Routing and Security Campaigns

Karl Olson, Bashayer Alharbi, Gregory Cusack, Eric Keller
University of Colorado Boulder, Boulder, USA
Email: {karl.olson, bashayer.alharbi, gregory.cusack, eric.keller}@colorado.edu

***Abstract*—Despite the desire to fix BGP underlying security shortcomings, present solutions, such as RPKI, have struggled to achieve broad adoption. Focusing on providers' needs first can incentivize the adoption of platforms that better enable the integration of security mechanisms, thereby overcoming stagnation and deployment barriers. To demonstrate this concept, we propose a real-time global routing database of network data that providers could leverage to support the management, troubleshooting, and business needs of their own networks. We show how broadly sharing information about networks, such as link usage and cost, can be leveraged to obtain *business optimal* routing decisions that could reduce provider transit costs by an average of 10% over standard BGP route selection or to provide route selection based on customer-defined security requirements– incentives for a business to adopt. We then leverage the same database to show how security solutions, similar to RPKI and BGPsec, could easily be adopted to further enhance internet security outcomes, demonstrating an incentivized approach to security adoption.**

***Index Terms*—BGP, networking, security, database, internet**

I. Introduction

Designed over thirty years ago, the Border Gateway Protocol (BGP) remains a necessary protocol for Internet routing. However, BGP has well-documented security challenges, including human error [1], [2], data falsification [1], [3], protocol manipulation [4], and data misuse [5]. A number of approaches have been proposed to strengthen BGP security, such as validating route origination [6], [7], ensuring path traversal [8], [9], and applying best practices to limit bad routes propagation [10], [11].

One common issue with these proposals is the lack of incentives for deployment. Let us use Resource Public Key Infrastructure (RPKI) as an example. RPKI provides route origin validation (ROV), which ensures that route advertisements originate from authorized ASes, and is considered the most widely deployed routing security solution. Yet, CloudFlare estimates that only 6.5% of ASNs are protected by ROV, covering just 5.6% of internet users worldwide [12]. While underdeveloped areas skew results, a recent large scale study [13] found many networks do not fully deploy ROV. They went further and interviewed 82 network operators to understand the reasons. Two main reasons were found: first, RPKI is minimally effective without broad participation and attackers can still circumvent it [14], [15]. Second, it takes a capital and operational investment to deploy RPKI, but does not bring additional business value (i.e., providers can not make more money, or spend less). From a business perspective, solutions that ease management, troubleshooting, or enable efficiencies serve as strong incentives for adoption. If we can align the components of a security solution to these incentives, we can increase the likelihood of adoption.

To demonstrate this, we present **DND-Db** (a Democratized Network Data Database), a real-time Internet routing database. DND-Db provides increased visibility for network administrators to manage, troubleshoot, and leverage network insights for business decisions– key incentives for a business. It goes beyond external measurement infrastructure (such as RouteViews [16]) as it allows direct participation by the network operators to incorporate information about their local environment (beyond routing). It also goes beyond what network operators may already be using for network management, as it decouples the information from a given commercial solution and can relate information to global data.

This solution seems both simple and impossible. Simple in what we propose is just a database of routing information. While true, we show the power of decoupling network data from the underlying network. Impossible in that it is proposing a central solution that sounds like it needs broad community participation. However, our proposal centers on the idea that broad adoption is ***not*** necessary for a solution to be beneficial if it provides value to individual providers.

Utilizing our prototype, we emulated a real-world architecture of 53 ASNs, leveraging Service Level Agreement and network transit cost data, stored in our global database, to enable *business optimal* topology outcomes. By including business requirements into the network decision making, we show provider savings of up to 10% on overall transit costs by leveraging broader types of data for route selection, demonstrating an incentive for providers adopting a global database. We further show a sliding scale of security, aligned to unique provider or customer requirements. This could potentially introduce new product or pricing tiers and be leveraged to select 100% compliant routes without requiring broad adoption, further incentivizing early adopters. Finally, we extrapolated the emulation results to outline initial parameters for deploying a global network database.

II. Demonstrating an Incentivized Approach

A. Aligning Solutions to Fit Business Needs First

Network providers are commercial entities whose primary product is providing connectivity at the lowest cost, turning a profit, or maintaining a competitive edge. Security, while

important, is not a formal requirement—delivery is. Here, we motivate our incentive based approach through the use of service level agreements as an example.

B. Service Level Agreements

Service Level Agreements (SLAs) define provider guarantees, customer expectations, and remedies for non-performance. In transit services, SLAs typically cover metrics like latency, packet delivery, and network availability [17].

Because SLAs are central to service providers, ensuring network performance is critical but must be balanced against costs to avoid over- or under-subscribing the network. To maintain performance and account for SLA impacts, administrators incorporate redundancy and balance loads for resilience. This approach has two problems: First, it creates a reactive network that responds only to a *network* event, potentially producing outcomes that are *network optimal* but *SLA detrimental*. For example, a planned failover link may become oversubscribed, causing an *SLA violation*. Second, it requires providers to pay for extra capacity that remains underutilized.

C. Leveraging DND-Db for an SLA Optimal Network

To prevent SLA violations, network and business data must be leveraged to make *SLA-optimal* network decisions that align with business objectives and tailor network responses. We demonstrate this with three scenarios where a provider could use a global network database to optimize outcomes while maintaining SLA requirements:

1) Sole Participant: To avoid slow deployment, a solution must provide value even if adopted by a single provider. A sole participant can use the global database to manage local network statistics, such as live measurements and historical trends, alongside business data like contractual performance measures to maintain SLA compliance. By storing historical route data, providers can identify congestion patterns and preemptively reroute traffic. If an alternate link is down (reported via DND-Db), adjustments can be made dynamically. While some routes have higher latency, they may still meet most SLAs. Leveraging DND-Db enables real-time traffic optimization, ensuring compliance without manual intervention.

2) Partial Participation: As participation grows and other ASes share data, we can optimize local traffic forwarding based on knowledge external to our AS. For example, if an AS reports that a specific link has exceeded the 90% capacity threshold and sends an "alert" to DND-Db. Subscribed providers can proactively forwarded traffic to alternate paths, avoiding SLA violation and potentially helping the impacted AS until the high usage subsides.

3) Full Participation: With broad participation and data sharing, new routing and optimization paradigms can be realized. If every node shares its transit costs and link performance metrics, we could design SLA-optimal routing while balancing additional metrics such as cost, capacity, or security. Different SLA performance thresholds can establish new pricing tiers while ensuring that both obligations and costs are maintained optimally, supporting a differing network objectives and customer requirements.

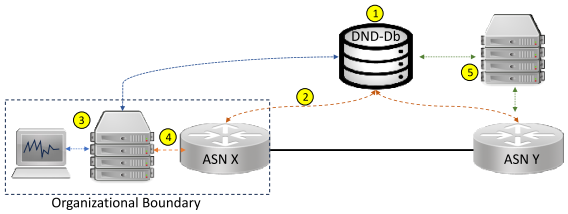

Fig. 1. **Overall Architecture.**

III. Architecture

A solution could take many forms to incentivize adoption, and we present one– a global database of real-time routing and infrastructure information. Importantly, this incorporates information beyond what can be measured by globally deployed route measurement infrastructure, and operates without humans in the loop so information does not go stale.

A. Overview

Figure 1 shows our architecture. A central routing database (label 1) provides processing, storage, and access-controlled APIs for read/write access. Each AS interacts via APIs for inline (label 2) (e.g., validation or inform routing decisions, similar to RPKI) or offline processing (label 3) (e.g., reporting, analysis, configuration, or management support). Locally cached data (label 4) can be used for performance-sensitive actions. Measurement infrastructure (label 5) could provide initial data to the database, benefiting early adopters and democratize network measurement data for broader use.

What to store: We consider three tiers of information. First is routing information as seen by each AS– significant insight about the Internet can be gleaned from this, as demonstrated by academic measurements [18], [19], which can be valuable in supporting local routing policy, security, or strategic investments. The second is application-specific information. For example, to support an RPKI-like solution, additional information about prefix ownership would be required. Finally, private data allows a network provider to link internal data, such as IP Address Management (IPAM), with routing information to verify route announcements. While this information could be maintained locally by the provider and solutions for this already exist [20], [21], decoupling and sharing network data can enable new opportunities for a more secure and dynamic internet.

B. Processing

While the database is logically centralized for storage, data processing occurs at each network provider and is categorized as: inline and offline.

Inline processing is logic at the router, enabling local routing decisions or real-time updates to the database. Two areas in the BGP control plane require additional processing by an AS (illustrated in Figure 2 as green boxes). The first are the inbound and outbound policy engines. The inbound policy engine verifies the prefix owner's record from the global database before accepting or rejecting a route. Similarly, the outbound engine checks the provider's private IPAM to prevent

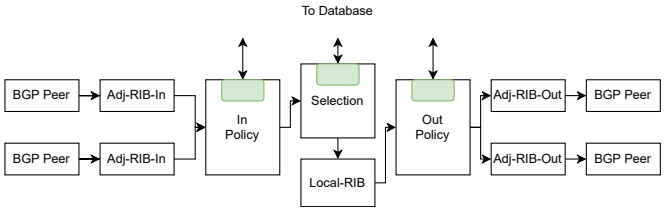

Fig. 2. **BGP Processing Pipeline.**

route conflicts or leaks. These engines also update the global database, e.g. a route advertisement by the AS would update the global database with new routing information, ensuring efficient management and data freshness.

The second processing point is within the BGP route selection process itself. For instance, paths can be preferred based on AS participation in security solutions, network stability, or routing through specific countries, enabling network providers to optimize routing for their own or their customers' requirements. Extensive inline processing might raise concerns, but risks can be mitigated through failover actions, caching, or new routing table designs. Since routers may not immediately support this additional processing, a proxy solution, as discussed in Section IV, can be used.

Offline processing complements inline processing and adds value for network management tasks or pre-establishing route configurations for anticipated network events. We see this as locally deployed software reading from the database for tasks like troubleshooting. While current commercial solutions already collect this data, our proposal decouples it from the local environment, enabling broader use and innovation.

IV. Prototype Implementation

We developed a prototype of DND-Db with two main components: a global database and a proxy to manage the interface between routers and the database.

A. Global Database Design

For our global database design, we chose MongoDB, allowing the customization of records for each AS. MongoDB's horizontal scaling and redundancy support larger processing scales without significant redesigns. It also includes security features to ensure data integrity and prevent unauthorized access [22]. Each AS maintains two types of records: **The public record** shares general AS state and relationships with others. **The private record** stores sensitive data such as private links, relationships, and non-public network performance. These records enable controlled information sharing while allowing providers to manage internal data securely.

B. Proxy Packet Handler

As a proof-of-concept implementation to handle the processing of packets and interaction with the database from each AS, we implemented an inline packet-processing proxy within each BIRD [23] router to intercept/process packets prior to reaching/leaving the router.

The proxy monitors incoming and outgoing packets and performs processing actions aligned with our goals. For example, to accept only routes where 50% or more of the ASNs along the path deploy RPKI, the proxy listens for BGP updates containing network layer reachability information (NLRI). Upon receiving an update, it identifies the ASes in the AS_Path and queries each AS's database record to verify RPKI deployment. A final RPKI participation score is calculated for the route. If the route does not meet a pre-defined threshold, it would be stripped from the BGP packet with the remainder forwarded to the router for processing. This approach allows control over route acceptance based on our defined requirements.

V. Evaluation

As DND-Db is a departure from traditional thinking for routing security, this section evaluates three core properties:

- **Value** (Section V-B): DND-Db incentivizes adoption by providing business value. We demonstrate how a global database can benefit administrators and businesses.
- **Equivalency (and beyond)** (Section V-C): We evaluate whether we can achieve equivalency with current state-of-the-art solutions in Internet security– RPKI for route origin validation, and BGPsec for path validation– while enabling opportunities for new paradigms.
- **Performance** (Section V-D): DND-Db must be globally scalable and not introduce overhead that noticeably impacts network convergence.

A. Experimental Setup

We emulated a 53 ASNs network using the SEED Internet Emulator [24], which leverages Docker containers running the BIRD routing daemon to replicate AS, allowing for a scalable network emulation. We used CAIDA AS-relationship data [25] to define real-world linkages and peering relationships. The environment included a central clique with six primary ASes, six Tier-1 ASes, eighteen Tier-2 ASes, and 23 customer-stub networks. These ASes were interconnected by 232 peer-to-peer and 76 provider-to-customer links. The resulting topology path lengths averaged 2.6, with a maximum of 5, approximating the real-world average of 5.3 [26]. Latency was incorporated by adjusting queuing discipline for select nodes relevant to our measurements. The environment used 196 Docker containers in total.

B. Results: Value

To drive adoption, DND-Db must offer business value. We evaluate two approaches that offer initial value: (i) SLA optimization of cost and (ii) route selection based on reported security deployments.

SLA Cost Optimized Route Selection. To illustrate SLA cost optimization, our database was initialized with publicly available transit costs. To optimize route selection based on cost, each router listens for route announcements and calculates the transit cost for each route using available AS_Path cost data. If a cheaper route is found, the proxy tagged it with a predefined BGP community value, prioritizing it over the existing route.

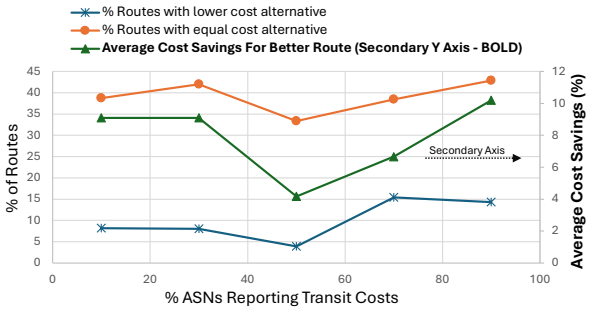

Fig. 3. **Availability of Cost Alternative SLA-Compatible Routes and Cost Savings Compared to Default BGP Route Selection.**

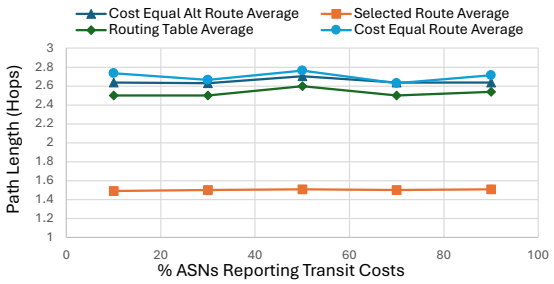

Fig. 4. **Average Hop Increase to Use a Lower Cost SLA-Compatible Route.**

To evaluate, we initialized non-peer-to-peer links according to the topology deployment rate (e.g., 10% represented 10% of peer-to-customer links reporting transit costs) using publicly available transit costs for 10 Gb/s links in North America [27]. Since transit costs can vary, we assumed 85% usage rate. Routing tables were dumped pre- and post-experiment to compare transit costs.

SLA Compatible Cost-optimal Route Selection Results. Based on the proxy routes data at the varying deployment rates (Figure 3), We identify opportunities where providers gain value. Compared to default BGP selection, up to 15% of routes were lower cost. 40% of non-selected routes were equal cost to the chosen route, allowing a provider to adjust local forwarding without additional costs. On average, achieving lower cost required one additional hop (Figure 4), with negligible latency increase, an important SLA consideration.

While a provider affect only local forwarding, broad participation would enable holistic topology optimizations. Even sole participants can adjust network actions to support business objectives locally, providing a valuable opportunity for managing business requirements. While upstream costs are not necessarily passed directly to downstream customers, they have an indirect effect, providing a metric to plan against.

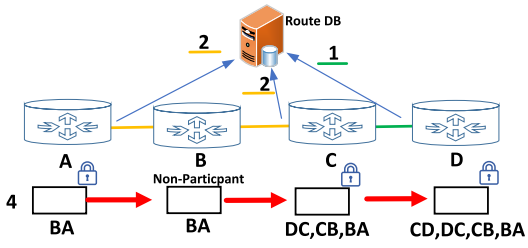

Fig. 5. **Establishing Path and Topology Validation.**

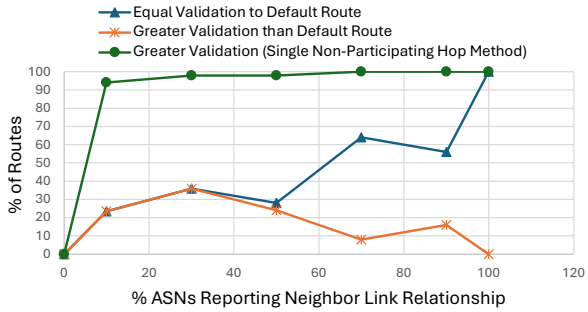

Fig. 6. **Route Selection Opportunity Based on % of Path Validated.** Using random neighbor link selection, we show the opportunity to select alternative routes based on path validation compared to the default BGP path. Fully validated routes with the single gap method are also demonstrated.

Validated Path Route Selection. To reference the existence of a path between two ASNs, each ASN publishes its BGP neighbor relationship to the database upon initialization (Figure 5). We defined validation levels based on whether one or both ASes attest to the link: First, each participating AS publishes its neighbor relationships to the database in its public record, allowing the distant neighbor to sign the same record as an *attestation*. This represents a partially validated link (label 2). The distant neighbor can attest to this record by signing the local AS's path utilizing the granted permissions. The path is then stored in the global database as a validated topology segment (label 1). The same process would occur for the distant ASN, resulting in two records of the path, one for each ASN's public record. Using the same approach in the SLA example, each proxy selects routes based on the percentage of a path validated, enabling route selection based on security metrics defined by the provider.

Validated Path Route Selection Results. To assess value, we compared the validated routes selected by the normal BGP process over a range of participation rates, as shown in Figure 6. At the lowest rates, both greater and equally validated route alternatives were available up to 30% of the time, diverging toward equal cost alternatives at higher participation rates. This is mostly due to BGP favoring the shortest path. Alternative routes typically incurred additional hops. The overall increase in path validation was minimal (less than 10%) compared to the selected BGP route. The high validation in this result was primarily due to the central position of our measurement node in the topology and the random assignment of published data along the primary path for this ASN's traffic.

The cost to select a route with better validation averaged 1-1.5 additional hops. However, if a provider offers better opportunity for a customer, this approach could incentivize businesses with new solutions. As an alternate, we assessed routes availability where a single non-participant separated two participating nodes, allowing for a "degree of assurance" metric to validate a path against, shown in Figure 5 as node "B". This approach ensures that the path provided and the node between participants exist, resulting in nearly fully validated path across all participation rates. Network time data could further verify no malicious routing actions occurred.

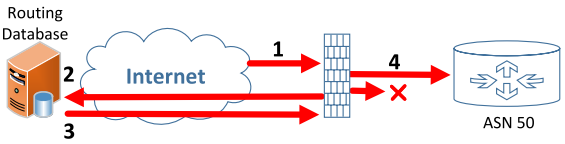

Fig. 7. **Route Origin Validation Using Routing Database.**

C. Results: Equivalency (and beyond)

Our goal is to increase adoption without compromising security. DND-Db must support state-of-the-art solutions for route origin validation (RPKI) and path validation (BGPsec). Previously, we presented how DND-Db could enable providers to select routes based on security metric. That showed a sliding scale of security, for example, a given path is 75% validated, whereas BGPsec is absolute. In this case, DND-Db could be configured to require 100% path validation, making it equivalent to BGPsec.

To demonstrate that our database equivalency to RPKI, we implemented a comparable approach, as shown in Figure 7. The proxy acted inline (label 1), listening for inbound advertisements. For each received advertisement, the proxy conducted a record lookup to the global database to validate ownership (label 2). If a validated record was found, the proxy correlated the received advertisement (label 3) with the data record to decide whether to accept the announcement (label 4). We randomly assigned this approach to a defined percentage of ASes within our topology before executing a prefix hijacking attack. To ensure a direct comparison, we deployed RPKI on the same ASes in subsequent tests. After the attack, we assessed each router's routing table for the hijacked route.

Demonstrating Equivalency to RPKI Results. We show that utilizing a database provides equivalency to RPKI while offering broader evolution. Across the varying deployment scales, we found complete agreement between both approaches, as shown in Figure 8. Measurements below the one-to-one line reflect random selection rounding reduction and not a drop in security, while measurements above reflect security extended to non-participants due to an upstream provider rejecting the bad route and preventing further propagation.

D. Results: Performance

Two key aspects of the architecture that introduce potential performance concerns. First is the reliance on a global logically central database, which must scale with adoption. Second is the extra overhead in processing routing messages, which must remain minimal to avoid delaying BGP convergence. We evaluate both and contextualize our results within global requirements. We find that DND-Db is scalable, practical, and can meet the needs of a global deployment.

1) Scaling of the central database: As database load scales with adoption, we used two approaches to assess topology performance. First, we established a single measurement node to record inbound BGP updates by restarting the BGP process for each router. After resets, we allowed updates to settle before triggering the next event. This allowed gathering metrics for the number of requests being generated to the database, along with additional metrics like average path length and

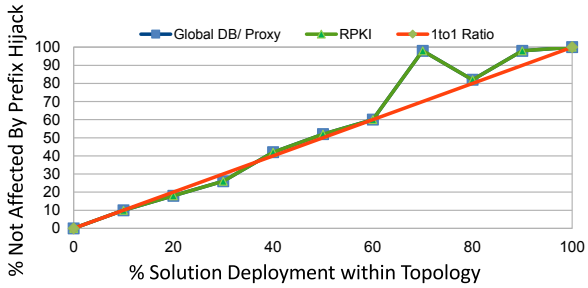

Fig. 8. **RPKI/Database Approach Comparison.** In a 1to1 comparison, our global DB approach was able to match RPKI results (hidden behind RPKI line) while offering greater flexibility to leverage data for other needs.

prefix counts per update, which we compare to global metrics measured from RIS [28].

Since this approach generated limited traffic rates before negatively affecting topology routing, our second approach used a custom traffic generator to simulate higher BGP update rates. This measured proxy performance and capacity across varying BGP update rates, which we compared to global rates.

TABLE I
SELECT TOPOLOGY AND BGP UPDATE METRICS.
METRICS CAPTURED THROUGH A SIMULATION OF OUTAGE EVENTS AT EACH ROUTER.

Simulated Topology BGP Update Metrics	
Statistic	**Total**
Unique ASNs / Routers	53 / 155
AS_Path Length (Max/Avg.)	5 / 2.63
BGP Updates Recv. (Total/Non-Wthdrwl)	977 / 757
Updates/Event (Total/Non-Wthdrwl)	6.30 / 4.88
DB Lookups (ROV/Path)	1532 / 1993
# Prefixs/Update (Max/Avg)	7 / 2.02
# Paths / Update	2.63
DB Req/Update (Max/Avg.) **Adj. For Total Rate	12 / 4.63

Topology Performance Results. For our topology generated BGP update approach, we obtained metrics from a complete measurement run, as shown in Table I. On average, 4.63 database lookups occurred per BGP update, slightly lower than our calculated global rate of 6.04 from [26], assuming both path and prefix validation were implemented. Utilizing these averages, we calculated the expected rate of server requests across varying update rates and compared to our results. This approach can be valuable for network simulations where representative traffic across the topology is desired. Rates of up to 3 updates/sec resulted from a router reset every two seconds without loss. Higher rates reached approximately eight updates before topology resets, significantly affecting packet delivery. We calculated peak throughput by generating server request packets at 20k requests/s, with steady-state operation averaging 11k requests/s.

2) Processing overhead at the router: Using our update traffic generator, we measured overall processing time and throughput capacity across varying network latencies, as shown in Figure 9. The steady state processing times were used to mathematically calculate the maximum steady state throughput and queuing capacity. We assumed a queue capacity of 1000 packets, the default size in most systems, to include our proxy. Results of the maximum throughput and queuing capacity are in Figure 10. Notably, network latency

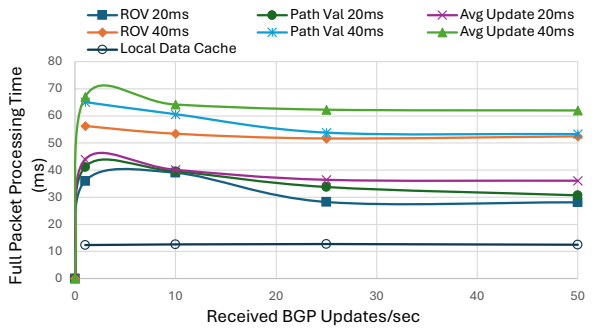

Fig. 9. **Latency Effect on Processing Time.** Latency effects on processing time for single (ROV) and multiple validations (Full Path) are shown. An average BGP update with one prefix and an AS_Path length of five calculates overall throughput.

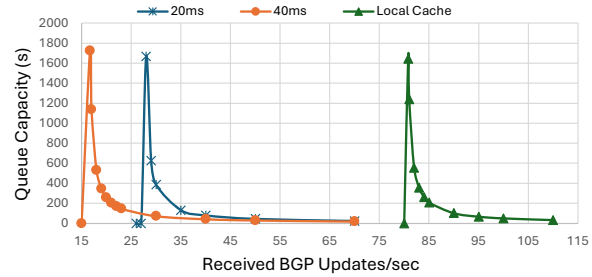

Fig. 10. **Overall BGP Update Processing Throughput and Packet Queue Capacity.** Peak throughput of 16.12, 27.4, and 80.19 BGP updates/s are demonstrated across select latency's. Queue capacity is based on a maximum queue of 1000 packets, demonstrating potential ability to scale for peak events.

significantly affects processing rates, demonstrating the need for careful consideration of database design and location(s).

3) Extrapolating to Global Requirements: We provide two sets of values: the first assesses requirements for a single request to the database, as used in a solution like ROV. The second approach calculates requirements based on multiple concurrent actions, such as with an AS_Path validation. These two approaches allow extrapolate requirements to other designs. For example, verifying average performance along a route need to request metrics from every node. If multiple metrics were desired per node, we would double the requirement. For calculations, we rely on the values in Table II and our previously calculated metrics for an average BGP update.

TABLE II
GLOBAL BGPUPDATE STATS AND SINGLE AS SAMPLING.
NOTE: ASN65000 VALUES ARE TABLE UPDATES FOR A 1-HOUR WINDOW, NOT BGP UPDATE RATES. * INDICATES A SIGNIFICANT ANOMALY; AVERAGES WITHOUT ANOMALIES ARE IN PARENTHESES.

BGP Global Stats [29]		AS65000 Stats [26]
BGP Statistic	**Global Value**	**ASN Value (max, avg)**
Active ASNs	75,042	
BGPUpdate MSGs/day	200,000 (2.31/s)	1,477,363 (17.1/s)
–AS_Path Change	140,000 (1.62/s)	42500* (1380), 300
–Next Hop AS Change	60,000 (.694/s)	1100, 200
–AS_Path Prepend Change	7,000 (.081/s)	91, 30
–Origin Change	2,000 (.023/s)	20,000* (2400), 270
AS_Path Length (Max / Avg)	- , 5	13, 5.37

VI. RELATED WORK

While our work is uniquely focused on establishing network provider value to incentivize adoption, the related work falls into two broad categories: routing databases and routing security.

Routing Databases. Existing routing databases and measurement infrastructure provide valuable network data, but correlating data across disparate systems remains a challenge. A centralized approach offers value, as we propose. Internet measurement methods support research, visualization, and troubleshooting [30]–[33]. However, these approaches work to reverse-engineer networks from limited vantage points. A stronger approach is to build infrastructure to enable broad community participation in reporting, sharing, and leveraging network data, which our work uniquely argues for.

Commercial solutions offering active monitoring and network metrics [20], [21], [34] are widely used but provide a narrow, limiting external value. While approaches such as remote access accounts extend visibility, they impose management costs and are not scalable for broad adoption. Recent proposals have explored blockchains for BGP security, including tamper resistant BGP messaging [35] and resource management [36]–[38], with functionality similar to RPKI. However, these approaches struggle with throughput necessary for internet scale, which remains an active area of research.

Routing Security. Over the past thirty years, number of BGP security solutions have been proposed, as captured in a several survey papers [39], [40]. These solutions often address only part of BGP's challenges while adding administrative costs and requiring broad adoption, discouraging early adopters. Current BGP security solutions, RPKI and BGPsec, suffer from these exact challenges. However, DND-Db differs by offering immediate value even with limited participation, making it more accessible to early adopters.

To ease the impact of another solution deployment, modifications to RPKI infrastructure could be leveraged to reduce deployment barriers while building on RPKI's progress. Regardless, for successful adoption, a solution must provide immediate value or offer incentives by addressing challenges providers face. By leveraging existing infrastructure, such as RPKI and network measurement initiatives, DND-Db enhances these solutions without increasing complexity, fostering broader adoption and long-term sustainability.

VII. CONCLUSIONS AND FUTURE WORK

Our proposal focuses on building solutions to better enable network administrators to perform their core functions. By aligning solutions with the administrator and their business, we can establish components for easier security mechanisms deployment, leading to broader adoption. To support this vision, network data must be more accessible and easier to leverage. We artificially constrain success by proposing solutions that are either narrowly scoped or static in design. Instead, we should pursue a more flexible and evolvable approach. We demonstrated one such approach, a global routing database, to enable this evolvable design. For future iterations, we plan to fully incorporate our design into BIRD, demonstrating seamless integration into modern routing software and extended for new capabilities. We plan to then leverage DND-Db in additional network operations use cases, demonstrating the synergy of a security with a network management solution.

REFERENCES

[1] Mahajan, R., Wetherall, D., & Anderson, T. (2002). Understanding BGP misconfiguration. ACM SIGCOMM Comp. Comm. Review, 32(4), 3-16.

[2] Cho, S., Fontugne, R., Cho, K., Dainotti, A., & Gill, P. (2019, June). BGP hijacking classification. In 2019 Network Traffic Measurement and Analysis Conference (TMA) (pp. 25-32). IEEE.

[3] Goldberg, S., Halevi, S., Jaggard, A. D., Ramachandran, V., & Wright, R. N. (2008). Rationality and traffic attraction: Incentives for honest path announcements in BGP. ACM SIGCOMM Data Communication

[4] Mitseva, A., Panchenko, A., & Engel, T. (2018). The state of affairs in BGP security: A survey of attacks and defenses. Computer Communications, 124, 45-60.

[5] Song, Y., Venkataramani, A., & Gao, L. (2016). Identifying and addressing reachability and policy attacks in "secure" BGP. IEEE/ACM Transactions on Networking, 24(5), 2969-2982.

[6] Lepinski, M., & Kent, S. (2012). An infrastructure to support secure internet routing (No. rfc6480).

[7] Wan, T., Kranakis, E., & van Oorschot, P. C. (2005, February). Pretty Secure BGP, psBGP. In NDSS.

[8] Kent, S., Lynn, C., & Seo, K. (2000). Secure border gateway protocol (S-BGP). IEEE Journal on Communications, 18(4), 582-592.

[9] Lepinski, M., & Sriram, K. (Eds.). (2017). RFC 8205: BGPsec Protocol Specification.

[10] Karlin, J., Forrest, S., & Rexford, J. (2006). Pretty good BGP: Improving BGP by cautiously adopting routes. IEEE International Conference on Network Protocols (pp. 290-299).

[11] Chen, E., & Rekhter, Y. (2008). Outbound route filtering capability for BGP-4 (No. rfc5291).

[12] Rodrigues, C. and Giotsas, V. (2022). RHelping build a safer Internet by measuring BGP RPKI Route Origin Validation. Available: https://blog.cloudflare.com/rpki-updates-dat

[13] Qin, L., Chen, L., Li, D., Ye, H., & Wang, Y. (2024). Understanding Route Origin Validation (ROV) Deployment in the Real World and Why MANRS Action 1 Is Not Followed. In Network and Distributed System Security Symposium (NDSS 2024).

[14] Zhang, M., Zhang, X., Barbee, J., Zhang, Y., & Lin, Z. (2023). SoK: Security of Cross-chain Bridges: Attack Surfaces, Defenses, and Open Problems. arXiv preprint arXiv:2312.12573.

[15] Hlavacek, T., Jeitner, P., Mirdita, D., Shulman, H., & Waidner, M. (2022). Stalloris:RPKI downgrade attack. In 31st USENIX Security Symposium (USENIX Security 22) (pp. 4455-4471)

[16] University of Oregon Route Views Project. *Route Views Archive Project*. Available: http://www.routeviews.org/routeviews/

[17] Cogent. Global Network SLA. Available: www.cogentco.com/files/docs/network/performance/global_sla.pdf

[18] Alfroy, T., Holterbach, T., & Pelsser, C. (2022, October). MVP: Measuring Internet routing from the most valuable points. In Proceedings of the 22nd ACM Internet Measurement Conference (pp. 770-771).

[19] Schlinker, B., Cunha, I., Chiu, Y. C., Sundaresan, S., & Katz-Bassett, E. (2019, October). Internet performance from facebook's edge. In Proceedings of the Internet Measurement Conference (pp. 179-194).

[20] Domotz. Available: https://www.domotz.com/

[21] ManageEngine OpUtils. Available: https://www.manageengine.com/allowbreakproducts/oputils/

[22] MongoDB. Available: https://www.mongodb.com/

[23] CZ.NIC, z.s.p.o. BIRD Internet Routing Daemon. Available: https://bird.network.cz/

[24] SEED Internet Emulator. Available: https://seedsecuritylabs.org/

[25] Center for Applied Data Analysis. AS-Relationships Dataset. Available: https://www.caida.org/catalog/datasets/as-relationships-geo/

[26] BGP Instability Report. Available: https://bgp.potaroo.net/as2.0/bgp-active.html

[27] Snijders, J., Hargrave, W., Rechthien, K., Stucchi, M., & Hoogsteder, P. IXP Cost Comparison. Available: http://peering.exposed

[28] RIPE Routing Information Service. Available: https://www.ripe.net/analyse/internet-measurements/

[29] Houston, G. APNIC - BGP in 2023. Available: https://blog.apnic.net/2024/01/10/bgp-in-2023-bgp-updates/

[30] Shavitt, Y., & Shir, E. (2005). DIMES: Let the Internet measure itself. ACM SIGCOMM Computer Communication Review, 35(5), 71-74

[31] Durairajan, R., Ghosh, S., Tang, X., Barford, P., & Eriksson, B. (2013, August). Internet atlas: a geographic database of the internet. In Proceedings of the 5th ACM workshop on HotPlanet (pp. 15-20).

[32] Staff, R. N. (2015). Ripe atlas: A global internet measurement network. Internet Protocol Journal, 18(3), 2-26.

[33] Salamatian, L., Arnold, T., Cunha, Í., Zhu, J., Zhang, Y., Katz-Bassett, E., & Calder, M. (2023). Who Squats IPv4 Addresses?. ACM SIGCOMM Computer Communication Review, 53(1), 48-72.

[34] ThousandEyes - Digital Experience Monitoring. Available: https://www.thousandeyes.com/

[35] Hari, A., & Lakshman, T. V. (2016, November). The internet blockchain: A distributed, tamper-resistant transaction framework. In Proceedings of the 15th ACM workshop on hot topics in networks (pp. 204-210).

[36] Xing, Q., Wang, B., & Wang, X. (2018). Bgpcoin: Blockchain-based bgp security solution. Symmetry, 10(9), 408.

[37] He, G., Su, W., Gao, S., Yue, J., & Das, S. K. (2020). ROAchain: Securing ROA with blockchain. IEEE Transactions on Network and Service Management, 18(2), 1690-1705.

[38] Paillisse, J., Ferriol, M., Garcia, E., Latif, H., Piris, C., Lopez, A., ... & Cabellos, A. (2018, July). IPchain: Securing IP prefix allocation and delegation with blockchain. In 2018 IEEE International Conference on Internet of Things (iThings) (pp. 1236-1243). IEEE.

[39] Mitseva, A., Panchenko, A., & Engel, T. (2018). BGP security: A survey of attacks and defenses. Computer Communications, 124, 45-60.

[40] Testart, C. (2018). Reviewing a Historical Internet Vulnerability: Why Isn't BGP More Secure and What Can We Do About it?. TPRC.