Where is the Incentive?
Rethinking Approaches to Security in Networks

by

Karl Olson

B.S., Mechanical Engineering, United States Military Academy, 2004

M.S., Environmental Engineering, University of Missouri - Rolla, 2008

M.S., Info. Tech. Mgmt., University of Maryland - University College, 2010

M.S., Electrical Engineering, University of Missouri - Columbia, 2013

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Computer Science

2024

Committee Members:

Eric Keller, Chair

Dr. James Curry

Dr. Eric Rozner

Dr. Yueqi Chen

Dr. Ravi Starzl

Olson, Karl (Ph.D., Computer Science)

Where is the Incentive?

Rethinking Approaches to Security in Networks

Thesis directed by Professor Eric Keller

Challenges with network security are as old as the Internet itself. While not the first incident, the 1988 Morris worm highlighted a need for more robust network security practices. Yet after more than thirty years of security research, proposals, and solutions, the Internet continues to remain broadly exposed to nefarious attacks at all levels of operation, from the underlying protocols that the internet operates on to the gateway hardware and software implementations found consumer home networks.

In order to understand why network security remains an unsolved problem, despite decades of effort, this thesis investigates the role of incentivization as a primary driving factor for the adoption, or lack thereof, of security solutions. To do so, we first conduct a retrospective assessment of consumer gateway security surrounding the role of network address translation (NAT), which we use to identify overarching trends, pitfalls, and missed opportunities for stronger security outcomes. We note that while a perimeter based security model afforded by NAT has never been a strong security approach, the simplicity, default-deny baseline behavior, and uniformity of design necessitated by address scarcity all served as strong incentives for deployment and use.

With the broad availability of addresses under IPv6, manufacturers are no longer bound by the default-deny design that NAT necessitated. Whether or not manufacturers are incentivized to continue offering a comparable default security baseline, and do so effectively, is unclear. To answer this question, we perform an assessment of IPv6 implementation found in ten consumer gateways. What we find is that many of the same security pitfalls surrounding NAT are being repeated, demonstrating that the need for security is not a strong incentive for actual implementation.

Based on this analysis, we then consider what an incentivized approach to encourage security development and adoption could look like. For this we shift focus from the home gateway environment

to the Border Gateway Protocol (BGP). While BGP security is fundamentally a different operating environment than consumer networks, the underlying need for strong incentivization to encourage adoption of stronger security approaches is a similar challenge.

To demonstrate this concept, we propose a global routing database that network providers could primarily leverage to support management and troubleshooting of their own networks. Utilizing this database, we demonstrate how broadly democratizing network data can be beneficial to a provider and their business objectives, providing an initial incentive for adoption. We then show how security approaches, similar to RPKI and BGPsec, could easily be adopted to our design. By leveraging the same systems, opportunities for new network paradigms can easily be created, allowing providers to leverage network data more broadly in the use of business objectives and routing security.

Dedication

To my family: Rose, Silas & Gage - Thank you for all your support and love. This was a team effort.

Acknowledgements

Thank you to Rose, for supporting me through everything and being my rock. You carried a heavy load throughout this in order to ensure my success. For that I am always grateful.

To Silas and Gage for giving me the opportunity to listen in on kindergarten and 3rd grade all over again while we all studied from home together during Covid. Always seek to challenge your worldview. New perspectives are found everywhere and will have a lasting impact on how you grow, even after 40.

To everyone back in Grantsburg, WI, who has played a role in my life, somehow you had a piece in all this.

To the research crew - Sepideh Goodarzy, Maziyar Nazari, Azzam Alsudais, Marcelo Abranches, Erika Hunhoff, Dwight Browne, Bashayer Ahlharbi, Jack Wampler, and Fan Shen. Covid limited our interactions, but you all still had a part in inspiring and offering your advice when we were together.

To everyone at the Army Cyber Institute - Ed Sobiesk, Stephen Hamilton, Todd Arnold, Paul Maxwell, Josh Bundt, Joe Catudal, Doug Healy, Doug Price, Richard Shmel, John Fuller, JC Fernandes, and Shane Kohtz. Thank you for all you do and the support over the last two years of finishing this out.

To Eric Keller, my advisor, mentor, and professor - A huge thank you. Your support was amazing throughout. I could not have asked for a better advisor in all of this. I cannot thank you enough for taking me on and getting me through.

Contents

Chapter

Tables

Table

Figures

Figure

Chapter 1

Introduction

Depending on where you look, the network security landscape has either changed significantly over the last decade or not at all. Significantly in that traditional approaches to securing networks, which have commonly placed broad trust in the operation of local devices while focusing security controls on the network perimeter, are no longer a sufficient strategy for holistic and effective network defense. Not at all in that the same security challenges present thirty years ago, such as with BGP route hijacking, still remain a common occurrence despite proposals and mechanisms to protect against these attacks [33, 97, 19, 190]. The result of these competing paradigms is a network environment that continues to be plagued by security challenges for both consumers and the internet as a whole. Despite decades of security research, exposures in the core systems that provide internet connectivity across the globe have counter-intuitively increased over time [132, 29].

1.0.1 Home Networks

Within the home network, a lack of participation by consumers in defining their own security environment enables an out-sized role for gateway manufacturers to define and establish baseline security requirements. While these manufacturers may establish controls that are appropriately secure, objectives to provide a simple operational configuration that works broadly for a wide range of consumers likely outweigh the need for strong initial security baselines. Further, market incentives to release new features, timed to market cycles, can often result in rushed, incomplete, or insecure systems [103], placing responsibility for security solely on the consumer to understand and implement. While manufacturers commonly provide

follow-up patching of vulnerabilities after release, the majority of consumers are rarely aware of the need to conduct updates to their systems, if they even know how [142]. The end result in most cases is the implementation of a gateway that is insecure out of the box, and likely remains that way over the life of the device.

In light of these challenges, we first consider how manufacturers have evolved approaches to security over time within the home network environment, which we use as a gauge for determining the effectiveness of present incentives for improving security outcomes. To do so, we performed a retrospective assessment of security vulnerabilities within home gateways related to the use of NAT and the surrounding mechanisms meant to ease consumer involvement in managing their local security policies. What we find is that although address scarcity served as a strong incentive for a universal gateway operational architecture based on address translation, offering both an initial security baseline and operational familiarity to the consumer, competing incentives to ease consumer involvement and abstract security away from the user have counterintuitively degraded the overall security posture of these gateways over time. Worse, we show the amount of exposures present to a consumer are increasing, demonstrating that the present incentive structure for security is limited at best.

Looking forward, the address scarcity incentive that drove the necessity for NAT and its familiar default-deny security baseline is no longer present under IPv6. How manufacturers may respond to this change in incentive is unclear. While a well configured stateful firewall could offer the same familiarity and security approach that consumers are accustomed to, complexity of addressing and broader availability for application make this a more challenging environment to precisely define clear security policies within. To answer how manufacturers are currently being incentivized to ensure security in this new landscape, we assessed a variety of home gateways, the effectiveness of their security policies, mechanisms for control, and overall gateway security posture under IPv6. What we find is an environment defined by ambiguity, security exposures, and unfamiliarity for the consumer, demonstrating that past considerations and understanding of security are not a strong incentive for future assurances.

1.0.2 Internet Routing

The protocols that run our networks are no better. The primary protocol that connects the broader internet together, the border gateway protocol (BGP), is a product of the early internet when security was less of a consideration for design. Little has changed in the protocol since, leaving service providers with broad exposure to routing attacks. These routing attacks can have a broad range of effects ranging from network outages [33] to interception and observation [13] of sensitive network traffic. More complex attacks have also used these exposures in BGP to control the processing of crypto currency, allowing for the theft of millions of dollars [93]. Solutions to mitigate these security weaknesses exist, but often add complexity and cost to providers hosting these networks - strong disincentives for adoption and deployment.

Resource Public Key Infrastructure (RPKI) is one such security method for BGP. With RPKI, route announcements are signed by an owner and published to a public RPKI database. Providers then host a local RPKI server to cache these public records, allowing their local BGP routers to easily query the available records any time a new route announcement is received. In this way, ASes can verify received route announcements have originated from a true owner while filtering out routes that conflict or do not match these authoritative records, providing a level of assurance against malicious or accidental route announcements. Despite RPKI providing a level of security for BGP, implementation has been slow and currently sits at a 40% adoption rate despite ten years of availability [130].

Part of the challenge preventing large scale adoption of RPKI is the unclear value presented under partial deployment [7]. Being the first to implement provides little overall value in terms of security while adding complexity and cost to a provider's network. Cloudflare, one of the world's largest content service providers acknowledged that implementing RPKI across all their networks would be prohibitively expensive, complex, and burdensome to implement as designed [52]. More importantly, BGP works - as long as all participants do the right thing. Most do, leaving little incentive to adopt solutions with unclear partial deployment value while taking on clear costs to implement.

Despite decades of research, proposals, and development of approaches to secure these network systems, the internet continues to be broadly exposed to attacks and malicious actions [33, 124]. We

argue that this security malaise is a result of unclear or mismatched incentivization for security, where competing objectives such as cost, time, complexity, and market dynamics are stronger drivers for design and operation of our networks and the devices that enable them. To change this paradigm, we need solutions which incentivize security adoption and provide clear value to the people likely to adopt them.

In order to demonstrate how an incentivized approach could work, we first consider the needs of a network provider and what a system would look like to help them perform their core business effectively and efficiently. By developing a solution that first helps a network administrator perform their core function we can incentivize adoption of mechanisms that we can later build security solutions off of. In order to demonstrate this incentivized approach, we propose a global routing database of routing information that providers can use to effectively manage and troubleshoot their own networks. We then show how security mechanisms, both equivalent to RPKI and new novel approaches, can be integrated easily, allowing for broader security opportunities and easier adoption. In this way, we incentivize the adoption of security and help to prevent many of the disincentives present in current designs.

1.0.3 Outline

The rest of this dissertation is organized as follows: Chapter 2 provides our initial assessment of the home router security landscape and the competing goals and objectives of parties in implementing security. Chapter 3 details how a changing landscape that focuses on individual device security has failed to incentivize adoption of stronger security controls from gateway manufacturers and considers how incentivization for security could be achieved in this unique market. Chapter 4 changes focus from the consumer home network to the broader internet and the security challenges it presents. Here we present an incentive-based design approach to security, which focuses on offering immediate value to a provider, even if they are the only one to implement. By incentivizing adoption for things a provider cares about, we can introduce mechanisms for security to build off of, potentially breaking the stagnation surrounding adoption of present approaches.

In each of the chapters mentioned above, we first identify how the chapter's effort helps us understand the role of incentivization by utilizing the chapter's focus effort. We then present the results

of that effort before drawing conclusions and ties to following chapters. This dissertation then wraps up with a conclusion and a discussion around future directions for research.

Chapter 2

Assessing the Role of Incentivization in Security Deployment Over Time - What have we Learned from NAT?

We begin this dissertation with a foundational assessment of the home router security landscape in order to understand how manufacturers have evolved there approach to security over time and to identify key trends or emergent incentives driving (or not) security implementation. To do so, we perform our analysis utilizing the perimeter security paradigm necessitated by network address translation (NAT) as a common lens shared across manufacturer implementations. We taxonomize the present NAT-centric model of consumer gateway security through a survey of over 300 common vulnerabilities and exposures(CVEs) surrounding NAT and hole punching protocols. From this survey, we identify a growth in security exposures over time as manufacturers have worked to abstract security away from the user through mechanisms meant to ease management and configuration of user networks. Underlying this growth in security exposures are clear market incentives that prioritize simplicity for users with plug-and-play baselines that often provide minimal security. Given the repeatable and remarkable occurrence of exposures found in our assessment, it is apparent that current self-governed approach is doing little to move the bar for security and that other incentives are required to change this outcome.

2.1 Introduction

Since the first formal proposal for a tiered address translation mechanism in 1992 as RFC 1335, the role played by network address translation (NAT) towards the meteoric expansion of the Internet cannot be understated. A 2006 study estimated that 70% of all devices accessing the Internet did so

from behind a NAT[1] gateway[22]. In the context of residential networks, that value jumps to nearly 95% [107]. Without the widespread deployment of NAT, the Internet could not support the 40 billion connected devices today [60].

Originally intended to overcome address exhaustion concerns, NAT quickly grew as a mechanism to increase privacy and security by masking internal network topologies and providing a default connection filtering mechanism [66, 50, 165, 197]. While NAT was not originally intended to function as a security solution, it is often *the only access control mechanism* protecting residential networks as it necessarily prevents unsolicited ingress traffic from reaching internal hosts. Studies have shown that around two-thirds of users deploy devices with default configuration settings [46, 181, 142] making the default-deny behavior afforded by NAT one of the most influential access control security mechanisms within residential networks.

With the ongoing push by internet service providers to deploy IPv6, the addressing constraints that gave rise to a familiar security and configuration baseline with NAT are no longer required. Instead, gateway manufacturers are left to decide between two very different operational contexts for IPv6 within consumer gateways: a familiar "closed model" approach where the gateway again serves a security boundary for access to the internal network, or a second "open model" approach that aligns with the intended end-to-end design of the Internet [151].

In light of the ambiguity presented by IPv6 operation, this work provides thirty-year retrospective assessment of the access control model presented by NAT and associated hole-punching security abstractions commonly used to manage gateway security policy. We follow this review with a systematic analysis on how these mechanisms meant to ease consumer involvement in home network security have traditionally failed in practice. To do so, we compile and assess over 300 associated vulnerabilities from the National Vulnerability Database (NVD) and Mitre Common Vulnerabilities and Exposures (CVE) listings in order to assess common vulnerability weaknesses, exposures, and trends. From this review, we contextualize the current consumer gateway access control security model and key operational lessons in order to better understand and define requirements for IPv6. We conclude by answering how manufacturers are, at present,

[1] For the remainder of this paper we refer to IPv4 NAT usage as "NAT". When referring to IPv6, we precede the term with the IP protocol, e.g. "IPv6 NAT".

approaching the open-ended design requirements surrounding IPv6 operation. To do so, we conduct an assessment of ten IPv6 gateway default security policies, controls, and device behaviors, which we use to contextualize the challenges and differences consumers are likely to face in deploying an IPv6 gateway.

In conducting this retrospective assessment, we find three recurrent themes that have an impact on present and future designs for consumer gateways and networks. First, we see a recurring failure, both with NAT and now with IPv6, where lack of specificity within formal documents pave the way for disparate interpretations by gateway developers, often at the cost of consumer awareness and security. Second, failures to assess security in light of new use cases often result in unintended exposures. For example, the hole-punching security abstractions meant to ease consumer configuration have commonly presented an overall increase in gateway security exposures resulting from incorrect implementations, use of outdated or vulnerable software packages, or insecure default configurations. These challenges continue with IPv6 as many of these abstractions are being directly converted from IPv4 packages while failing to account for differences in operation and addressing present under IPv6. Third, the ability of a consumer to rely on the presence of a default deny stateful filtering policy is no longer assured. In many of the gateways we reviewed, not only is the consumer network broadly exposed under default IPv6 security policies, these exposures also require active involvement by the consumer to correct. This is a paradigm shift in expectation which goes counter to the demonstrated behavior of users to change device default configurations at present. At best, we can define IPv6 operation in consumer gateways as a "default expose" security posture.

The remainder of this chapter is organized as follows: In Section 2.2, we define the common security properties, operational models, supporting parties and attacker goals in order to provide a common understanding of the complex interrelationships involved in defining a common access control model for consumer gateways. In Section 2.3, we contextualize NAT as an access control mechanism within the networking stack to demonstrate the importance of the network layer access control boundary in consumer networks. We then survey and document the operational methods within NAT and hole-punching methods, highlighting the broad complexity and nuance operating within consumer gateways in Sections 2.4 and 2.5. We use this context to taxonomize the operational failures of both NAT and hole-punching methods in Section 2.6 and conclude with a trend analysis to show that consumer gateway

security has never been great in 2.7. A review of related work concludes our survey in 2.9.

2.2 Background - Consumer Gateway Security Models, Properties, and Stakeholders

The focal point of every consumer network, a gateway serves as the interconnect between the local, customer managed, network and the broader Internet. This out-sized role demands a balance between often competing objectives of security, configurability, and ease of operation for the consumer. In order to understand the challenges with maintaining this delicate balance, and to systematically assess outcomes where these objectives have failed in practice, we present a short review of competing gateway security models, operational properties, and identification of parties involved in establishing a gateway's overall security. We further define security from an adversarial perspective, identifying key objectives and goals an attacker may pursue in attempting to overcome gateway security measures.

2.2.1 Home Network Security Models

Consumer network security is commonly defined by the security model employed at the customer demarcation or edge. Here, a transition from the globally routable network backbone, typically managed by an ISP, to the internal or customer managed network occurs. The type of security model employed is commonly dictated by the default configuration employed by gateway manufacturers. We describe these default behaviors a consumer may experience below.

Closed Model - A perimeter defense approach that focuses security controls at the network edge to prevent access to an internal or trusted portion of the network. Here, security is primarily focused on preventing broad *network* access. Devices within the security boundary are generally free to communicate with each other absent more refined security measures such as virtual LANs, separate SSIDs or host-based filtering strategies.

Open Model - Communication in the open model strives for end-to-end reachability without need for address translation or arbitrary borders and restrictions. Responsibility for security is shifted away from the network perimeter to each connected *device*. This open model approach is commonly found with

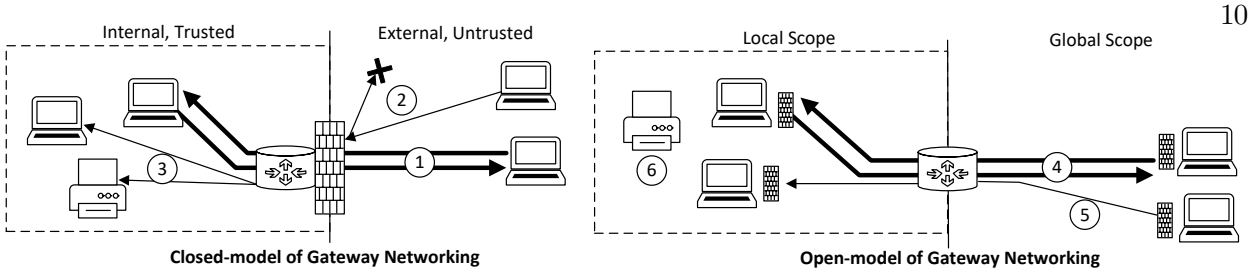

Figure 2.1: **Two Architectural Approaches to Gateway Networking.** In the closed model (left image), a gateway acts as the primary device to provide access control into a local network. 1) Network communications from the internal network are allowed outbound with connection state maintained to match and allow return traffic. 2) Unsolicited traffic is filtered at the network edge. 3) Internal devices are allowed to communicate freely absent other control mechanisms. In the open model (right image), network communications are end-to-end. Both 4) Outbound/return traffic and 5) inbound solicitation are allowed, shifting access control to the end devices. 6) Devices with no organic security mechanism may be exposed to the broader Internet.

early IP networks, when the scale and scope of the Internet was much smaller, and within the growing use of IPv6 networks where address space allows for the unique addressing of each connected device.

Hybrid Model - A layered approach to security that provides both perimeter security controls in conjunction with globally routable addressing for consumer devices. A hybrid model may take many forms, such as a network edge firewall with individual device policies, or through the re-implementation of address translation mechanisms similar to NAT.

2.2.2 Network Gateway Security Properties

While the aforementioned security models address the competing paradigms to gateway operation within a consumer network, security properties are the universal standards by which any device, protocol, or architecture should adhere. We briefly define these core security properties in order to establish a baseline for expected gateway security behavior.

Confidentiality is a property that ensures information is not disclosed to unauthorized individuals. In a secure gateway, at no time should information be leaked about the network, systems, or data to unauthorized parties. This premise assumes that gateways are established with secure default configurations, even though this may not occur in practice [118].

Integrity is the ability to guarantee system operation or data transmission remain true to their original trusted form or settings. Challenging this assurance is the fact that each and every component

making up a system must follow this principle in order to achieve a level of assurance for the whole device.

Availability guarantees that with all control mechanisms and security procedures in place, authorized individuals who require service are able to obtain such. In addition, a system operating in a secure manner should continue to operate and maintain individual services in the event of a component failure or compromise, as long as the failure does not introduce new vectors which could further system exposure.

Reliability, Authenticity, Non-Repudiation are recommended extensions of the CIA triad by ISO 2700 which further define how security goals may be achieved [80]. Authenticity guarantees that a user or system is who they say they are, often verified through a proof of validating credentials or through demonstration of specific knowledge, token, or fingerprint prior to access or communication. Non-Repudiation provides evidence or proof of actions which affect a system or data. This commonly occurs through system or event logging, such as through a security information and event management (SIEM) system. Finally, reliability concerns both the repeated and expected operation of a device for each action or transaction and the ability of a system to operate within the scope of expectation given an event.

2.2.3 Parties Involved

Security of a gateway is neither solely a manufacturer responsibility or a consumer task. It is a shared responsibility spread across many parties. Below we list the common parties, each of whom play a unique role in establishing the security of a consumer gateway, and by proxy, a consumer's network.

Consumers are network participants who are responsible for the local network and devices within it. This includes responsibility for the network gateway and any security policies they may chose to implement.

Developers/Manufacturers define and implement the components necessary to provide network and security services. Despite not having a direct role in the operation of a consumer's network, this group maintains an out-sized role in consumer network security due to implementation of default security settings, device patching, and inclusion (or absence) of security control mechanisms.

Internet Service Providers (ISPs) provide network service that connects users to the Internet, providing a consumer either an IPv4 gateway address or an IPv6 subnet via prefix delegation. While an ISP typically plays very little role in the security of a consumer's network, decisions to deploy and transition

to IPv6 can potentially have a profound impact on access control, which we discuss further in 2.8.

Standards Organizations define the operational requirements, considerations, and characteristics of functions used to provide network and security services. This allows developers to implement systems in a common and inter-operable way. On the other hand, vague definitions or open-ended requirements can present uncertainty and serve to hinder broader intents. Organizations such as the Internet Engineering Task Force (IETF), WiFi Alliance, and Open Connectivity Foundation commonly provide many of these standards present in home gateways.

2.2.4 Attacker Goals

Finally, in order to holistically assess access control in consumer gateways, we must consider the overall goals of an attack. We briefly define these attacker goals in order help frame the impacts security flaws may present. These categorizations align with prior works based on network attack goal classification in [89, 188, 102]. This is not intended to be a comprehensive review of attacker methodologies, but a common frame of reference from which to assess how NAT and associated hole-punching protocol failures have furthered attacker objectives in practice.

Access is when an attacker obtains the ability to utilize a system for their benefit. Access does not immediately imply administrative control and may be limited to solely viewing or monitoring of configuration settings and/or traffic.

Elevation is when an attacker gains the privilege to conduct actions or view information typically excluded from unprivileged users. With elevation comes the ability to perform additional actions to further individual goals.

Modification typically occurs when an attacker necessitates a change in system or data state to further ones objectives. For systems, this could be through assigning increased privileges, deactivating components, or other similar methods. With data, the contents of communication are modified such that the end result is a benefit provided to the attacker.

Denial of Service is the removal of a system's availability to provide ongoing service. This could be temporary in nature where service is restored upon conclusion of an attack, or it could be permanent

through means like physical destruction.

Information Gathering are the methods and techniques which enable an attacker to glean information to further objectives or goals. This information could come from unsecured communications, publicly available information, or through probing attacks.

2.2.5 Competing Goals and Security Trade-offs

Taken together, competing goals between stakeholders highlight the challenge of providing a secure yet functional consumer gateway. This complex security interrelationship poses a number of challenges to the consumer in particular. First, in order to play an active role in the security of their gateway, a consumer must have a working understanding of how a configuration settings, services, or applications tie to a defined security objective they seek to achieve. Second, they must have the ability to implement their action precisely (both in terms of operator skill and through an available security control mechanism) without further exposing their system or network. Stated more precisely, *a consumer's ability to precisely manage access control within their network is limited at best.* We highlight the mechanisms by which a consumer can enact access control measures to manage their security posture in the following section.

2.3 Residential Access Control Methods

Consumer gateways aim to be as simple as possible, requiring minimal consumer involvement beyond establishing a Service Set Identifier (SSID), Wi-Fi Protected Access (WPA) password, and any ISP-specific settings (such as a PPPoE username/password) [32, 30]. Beyond these initial configurations, a suite of protocols provide simple, often automatic, setup for connected devices and traffic flows such that the user does not interact with or receive feedback from the network unless a problem emerges [55, 192]. A default-deny security behavior enabled by NAT further provides a default security control to unsolicited inbound network traffic. As a result, operators need minimal technical understanding to establish and maintain a home network.

Within the local network a default permit security policy is commonly in effect, allowing connected devices to both freely communicate with each other and with external systems. Under this permissive

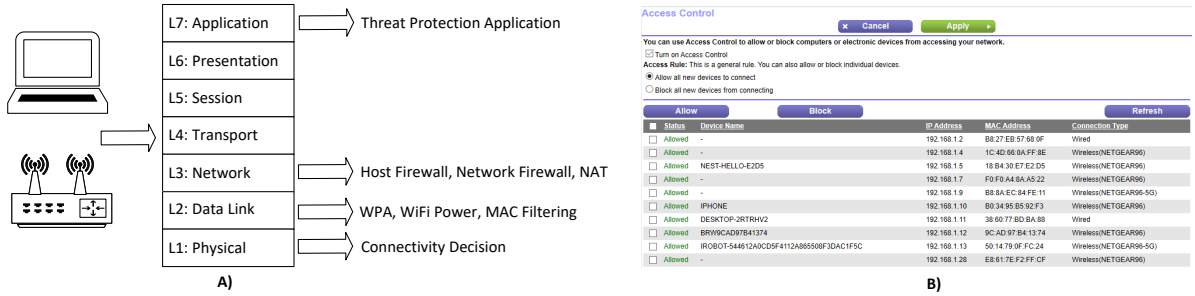

Figure 2.2: **Home Environment Access Control Methods.** A) An OSI layer view of typical access control methods available to consumers. B) Home systems and device manufacturers do little to ease identification of systems. While some devices show enough information to identify, many do not, making it hard to assess devices on a network for most users.

policy, consumers are left to presume that their connected devices do not behave maliciously, though evidence shows this assumption to consistently fail in practice [49, 182, 137].

Limiting specific device behaviors in this permissive environment is challenging at best. The simplicity and highly heterogeneous nature of consumer gateways has abstracted security away from the user, inhibiting the deployment of stronger access control measures and limiting mechanisms to precisely refine security policy. In particular, the options available to consumers to perform access control are limited; we show these available methods for each layer of the Open Systems Interconnection(OSI) model in Figure 2.2 and discuss in detail below.

2.3.1 Layer 1: Physical

Wired networks provide a simple, coarse, and effective access control mechanism: either a cable is connected to a network or it is not. This provides the user with a binary choice and is revocable without deep technical knowledge about the underlying system.

Wireless networks, however, suffer from problems which complicate low-level access control. The nature of RF transmissions in the 2.4GHz and 5GHz bands means that they frequently leak beyond the bounds of the physical location of the transmitter [148]. An adversary in an off-site location can collect these signals, disrupt, or attempt to connect to the network. While beam-forming [126] and secure arrays [196] can alleviate these issues, the user must still monitor the network for unauthorized devices.

If unauthorized devices are found on the wireless network, the options for remediation remain limited. No physical layer controls exist for evicting connected wireless devices, forcing the user to rely on weak controls at the data link layer.

2.3.2 Layer 2: Data Link

At layer 2, users can create a media access control (MAC) address filter to allow or block-list a known set of addresses, a feature typically disabled by default [164, 101]. The effectiveness of this control is limited; MAC addresses are a poor proxy for identity due to the simplicity of spoofing attacks (where an adversary attempts to bypass a allow-list or block-list by modifying a station's MAC address). Although some heuristic-based approaches exist to detect spoofing [158, 65], we consider these to be anomaly detection mechanisms and not access control policies.

Furthermore, some devices are capable of presenting multiple interfaces and MAC addresses (e.g., virtual machines with bridge networking and pass-through VoIP phones), which can further frustrate efforts to identify devices. Figure 2.2B demonstrates the vagueness of device identification commonly presented to a user managing a home gateway at layer 2.

With wireless, a user can restrict network access through mechanisms defined within the wireless encryption standards [8]. In the Wi-Fi Protected Access (WPA) scheme, for example, a Pre-Shared Key (PSK) is derived from a password, which is used to authenticate the device to the network. Password-based schemes provide a share-able mechanism for permitting access to a network. However, poor password choices, such as relying on dictionary words, family names, or even default manufacturer values [4] can allow adversaries to bypass this control. Control of these passwords are also often shared among family members or guests, increasing a user's exposure if a password is reused to access other systems [173]; some platforms (e.g., iOS and Windows) provide features which allow user to automatically share a wireless password with a nearby contact [14]. Once shared, these passwords are not easily revocable and the user must change the password and reconfigure all allowed devices.

As with physical layer controls, data link access controls are coarse. These typically apply to a single physical device and permit all traffic from the device once these controls are passed.

2.3.3 Layer 3-4: Network and Transport

These layers provide high granularity for access control with respect to individual traffic flows, both inside and outside the private network. The implementation of a stateful firewall initially seems ideal; such a system would allow the user to control both ingress and egress traffic through refined policy definition. However, firewalls require a detailed understanding of IP networking and the device or software responsible for managing the policy. These are difficult to implement correctly even for experts [194] and it is unlikely that the average user has/should have the requisite skills to configure a firewall.

In most residential IPv4 networks, a firewall provides marginal value due to the ubiquitous nature of address translation. While NAT was not originally designed to be a security feature, *it is occasionally the only ingress access control deployed on a home network [142].* The popular traditional NAT-PT mode of NAT (described in Section 2.4.0.1) effectively provides a security policy that prevents unsolicited inbound traffic from reaching the local network. This "security-through-unreachability" masks all devices behind the router providing a default privacy and security perimeter with little to no overhead effort for home network operators.

In contrast to the security provided by the default-deny policy of NAT, the broadly accepted and deployed permit policy for outbound traffic *assists* users in degrading their own security. Devices, such as TVs or IoT, commonly leverage this broadly permitted outbound traffic request to enable two-way communication with an external 3rd party, often unbeknownst to the user [1, 57]. Restricting this permissive outbound behavior is challenging at best for reasons previously mentioned.

2.3.4 Layers 5-7: Application / Host Based Security

At the highest layers of the OSI model users are again afforded with high granularity for access control on a *per-device* basis. Here, inclusion of host-based firewalls and automated policy mechanisms, such as an intrusion detection system, provide users a feature rich policy refinement platform. Ideally, this level of refinement and automation would be a boon for consumer security. In practice, there are many opportunities for failure.

First, detailed policy refinement again assumes an advanced level of knowledge, requiring an understanding of both networking and host policy metrics. Second, automation of policy creation using IDSes or similar methodologies provide an opaque level of security commensurate to a user's ability to ensure both timely and continued maintenance. Lastly, mechanisms by which a user may enforce policy at the host level are not universal. Competing objectives to provide users both the ability for detailed policy refinement and simple mechanisms by which to do it are often at odds with developers to provide timely and cost effective solutions. IoT or Smart Home devices are likely to forgo host based security altogether, leaving a consumer to either guess on the defensive posture organic to the system or rely on accurately implementing lower level controls [176].

2.3.5 User Considerations

In reviewing these access mechanisms, we see two clear takeaways: 1) fine-grained access control and the mechanisms by which to implement them require *some* level of knowledge and familiarity, and 2) we cannot assume that a user inherently has this level of knowledge or desire to implement such policy. Therefore, security in a consumer premise is commonly defined by the default security configuration and use of supporting mechanisms to automate policy on behalf of a consumer. This position appears to be supported by a number of studies which show that users rarely involve themselves with changing default configuration settings or do so in a way that improves their security [155, 85].

In the case of consumer home networks, the use of NAT and hole-punching mechanisms have commonly provided this default security policy and automation. With IPv6, this same common security baseline across gateway manufacturers is no longer required due to the broad availability of routable address space which no longer necessitates the use of NAT. In order to better define and understand what this transition means for consumer security moving forward, we believe it prudent to conduct a systematic review and assessment of NAT and associated hole-punching methods in order to glean lessons for IPv6 deployment.

2.4 NAT Operational Methods and Disparate Interpretations

The expectation for NAT to be a short-lived solution resulted in little guidance by the IETF on precise operational characteristics required [197]. This ambiguity lead to broad interpretations of NAT behavior by gateway manufacturers who were rushing to fill an explosive demand for consumer network connectivity. In the following section, we present a review of these diverse NAT operational methods and behaviors to highlight both the challenge and complex operating environment arising from ambiguity in specifications. While not every operational architecture is found within a home gateway, we include many of these to provide a complete view of the wide array of NAT methods employed in practice.

2.4.0.1 Traditional NAT and NAPT

Traditional NAT (NAT) maintains a single external IP address which is shared amongst all internal hosts. Sessions are uni-directional, meaning hosts from the internal network are able to establish a connection to the external network via a one-to-one address translation. Connection state is maintained within a forwarding table, allowing the NAT device to match inbound communications with the paired internal host as shown in Figure 2.3. At larger scales, a single external address limits the number of hosts that can request translation, resulting in two minor modifications commonly found in enterprise and consumer implementations: (1) Basic NAT, which maintains a pool of external addresses for sharing on a first-come-first-served basis and (2) Network Address Port Translation (NAPT or NAT-PT) which allows the multiplexing of many hosts into a single address through unique port assignments [169].

2.4.0.2 Bi-directional or Two-Way NAT

NAT relies on tracking a connection state to match return traffic to the correct internal host. For connections originating from the external network, there is no matching state. Further, the internal device may utilize a private address which are not routable in the global network. Two-way NAT enables inbound connection requests, as show in Figure 2.4. Here, external hosts may query a DNS server for the servicing gateway's external IP address. When an inbound request is received, the NAT gateway performs an address

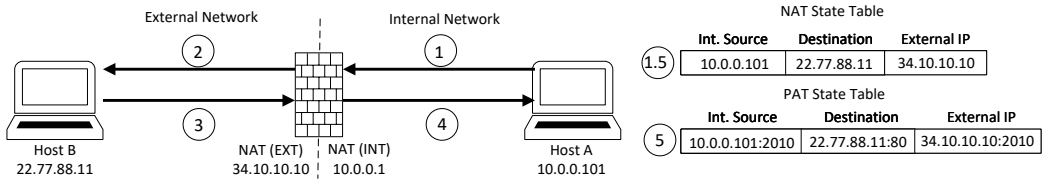

Figure 2.3: **Traditional NAT.** 1) Host A initiates connection to host B reaching a NAT gateway. 2) NAT gateway maps Host A IP address to external globally routable address, updates the IP packet to reflect the external interface IP and forwards packet to Host B. A connection state table within the gateway is updated to match return communications. 3) Host B responds using external global IP address of NAT gateway. 4 & 5) NAT gateway receives return packet, checks state table for matching internal host, updates destination address to reflect Host A and forwards packet. 6) With NAPT, connection state table maintains port assignment information to help support multiplexing of multiple clients sharing a single external IP.

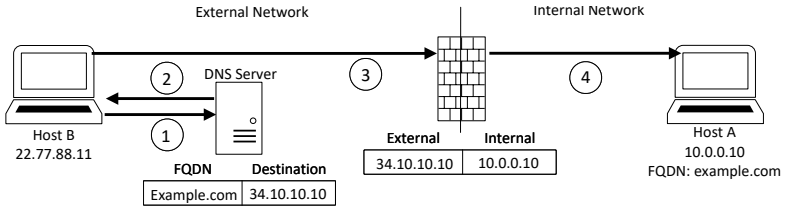

Figure 2.4: **Two-way NAT.** 1) Host B seeks to establish communication with Host A, located behind a NAT gateway, by first querying the public DNS server for the FQDN and external IP address of a hosted service. 2) DNS server responds with the public IP associated with FQDN. 3) Host B sends request to public interface of the NAT gateway, which checks the forwarding table for a static address mapping. 4) Request is forwarded to internal Host A.

search within the forwarding table, pairing the request with the internal matching host and forwarding the packet. This translation can be further defined by service, allowing gateways to host multiple applications or systems based on a listening port. Here, it is critical that the fully qualified domain names are end-to-end unique to avoid conflict in lookup and translation between external and internal hosts [169].

2.4.0.3 Twice NAT

With Twice-NAT, both the source and destination address of a packet are translated, as shown in Figure 2.5. This is desirable for a number of reasons. A company may not wish to update IP addressing after moving service providers resulting in overlapping public addresses; they may wish to rebind a request and redirect to another server; or they may have received a block of conflicting addresses from a merger or similar acquisition [169]. The concern is that an internal host may have the same routable address as an external host. When communication is executed internally, the request will not make it

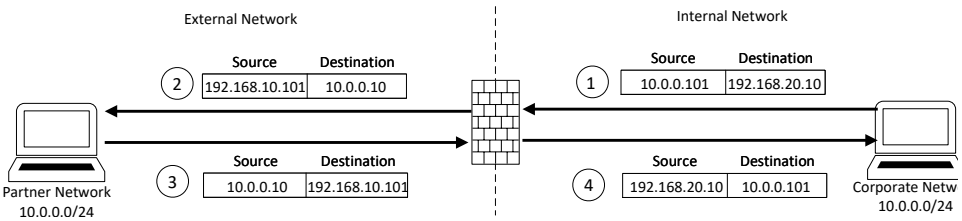

Figure 2.5: **Twice NAT.** 1) Host on corporate network seeks to communicate with partner office where both locations operate the same private IP space. Internal to the corporate network, the external partner network is assigned as an alternate IP space to prevent internal routing conflicts. 2) When communication from a corporate host reaches NAT gateway, pre-established translation rules update *BOTH* the source and destination packet to comply with routing and response on partner network. 3) Partner host responds to request forwarded by NAT device. 4) NAT device receives response from partner network and again translates BOTH source and destination IPs to route to requesting host on corporate network.

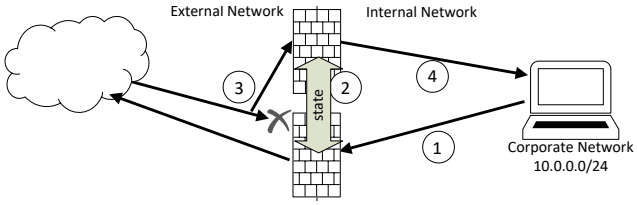

Figure 2.6: **Multi-homed NAT.** 1) Host on corporate network utilizes NAT to reach an external system. 2) NAT gateway updates state table and synchronizes state across all gateways. 3) In the event of an outage involving the primary NAT gateway, return traffic defaults to secondary gateway. 4) The secondary gateway finds the synchronized mapping in state table and forwards traffic to appropriate host on internal network.

to the external destination without translation. Likewise, a return request would have the same conflict. To overcome this, Twice NAT translates both the source and destination, keeping the proper routing path for internal and external hosts to communicate.

2.4.0.4 Multi-Homed NAT

One problem with NAT is that all communication must flow through the NAT gateway, making it a single point of failure in network architectures. To overcome this, Multi-homed NAT shares connection state information across multiple gateways, allowing a secondary gateway to transparently continue a session in the event the first gateway fails. Figure 2.6 demonstrates a typical configuration in multi-homed NAT networks. Here gateway #1 may be the primary NAT path which shares state information with gateway #2. In case of a failure all traffic is rerouted to gateway #2 transparently, ensuring communication is uninterrupted.

2.4.1 NAT Forwarding and Response Characteristics

In addition to the NAT architectures defined in RFC 2663, the development of the STUN protocol in RFC 3489 further defined the methodology and operation of NAT based based on forwarding and response characteristics employed by gateway manufacturers [147]:

2.4.1.1 Full-cone NAT

Full-cone NAT maps an internal host address ($IP_{Host}:Port_{Host}$) to an external gateway address ($extIP_{Gateway}:extPort_{Gateway}$). Any communication sent from an internal host will be translated by the gateway to the external address prior to forwarding to the target destination. Any external/return communication sent to the gateway's external interface ($extIP_{Gateway}:extPort_{Gateway}$) will in turn be translated and forwarded to the internal host ($IP_{Host}:Port_{Host}$), regardless of which external host is trying to communicate. In one 2008 study of NAT behavior deployments, full-cone NAT occurred in an estimated 37% of consumer gateway implementations [122].

2.4.1.2 Address-restricted cone NAT

With address-restricted cone NAT, the mapping and communication process is the same as full-cone. However, with address-restricted cone a state table is maintained to track communications and only the specific external host ($IP_{Ext.Host}:Port_{Ext.Host}$) may traverse the gateway on return. Ports do not play a role other than for the translation mapping in the NAT device. Therefore any port ($IP_{Ext.Host}:Port_{Any}$) may communicate with the internal host ($IP_{Host}:Port_{Host}$) upon a return response. Despite the increase in security afforded by restricting external hosts, address-restricted NAT was found in less than 5% of residential gateways [122].

2.4.1.3 Port-restricted cone NAT

Port-restricted cone NAT further limits operation of Address-restricted cone NAT. Here, an external host ($IP_{Ext.Host}:Port_{Ext.Host}$) can send packets to an internal host ($IP_{Host}:Port_{Host}$) only if the internal host has previously sent a packet to $IP_{Ext.Host}:Port_{Ext.Host}$. This methodology restricts communication

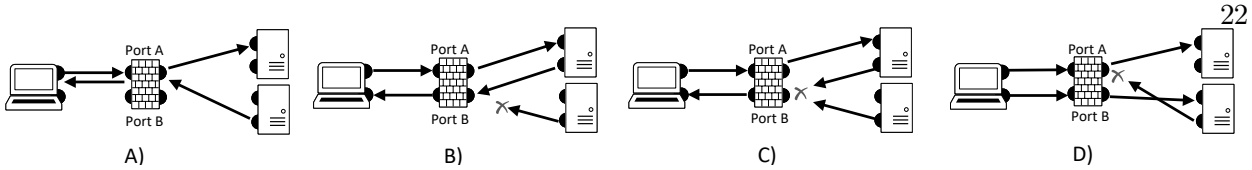

Figure 2.7: **NAT Operational Methods.** A) Full-cone NAT B) Address-restricted Cone NAT C) Port-restricted Cone NAT D) Symmetric NAT

in the forwarding table by both IP and port. Both port and address restricted NAT methods comprise the most common method of NAT implementation in consumer gateways representing nearly 51% of all devices [122].

2.4.1.4 Symmetric NAT

Each request from the same internal IP address and port ($IP_{Host}:Port_{Host}$) to a specific destination IP address and port ($IP_{Ext.Host}:Port_{Ext.Host}$) is mapped to a unique external gateway source IP address and port ($extIP_{Gateway_{U}niq.}:extPort_{Gateway_{U}niq.}$). If the internal host then sends a packet with the same source address and port but to a different destination, a new mapping is established in the translation table. Only an external host at $IP_{Ext.Host}:Port_{Ext.Host}$ that receives a packet from an internal host can send a return packet using $IP_{Ext.Host}:Port_{Ext.Host}$. Symmetric NAT is the least common comprising less than 5% of all consumer gateway implementations[122], despite presenting the strongest assurance for access control.

2.4.2 Proprietary Vendor Implementations

Further challenging the recognition of a single defined operation for NAT are behaviors often unique to a specific vendor implementation. These device specific behaviors provide unique or varying response characteristics and commonly include areas such as port selection methods, TCP state tracking, filtering response behaviors, timer defaults, and sequencing preservation approaches, to name a few [63, 82, 3]. To use the port selection as an example, some gateways select ports sequentially for use, another gateway may randomize port selection, and even another may sequentially check if any ports were recently closed for reuse before trying another approach [63].

Often these response characteristics are undocumented, requiring a consumer to conduct detailed

testing of their gateway in order to fully understand their device's operation. While we note this is a very untenable and far-fetched proposal, understanding these nuanced aspects do play an indirect role in router security through the need to potentially introduce or operate multiple hole-punching methods which address the many use cases [120]. This in turn increases a consumer's overall exposure, requiring assured implementation of additional protocols to guarantee a gateway's overall security. As we highlight in Section 2.7, this is rarely achieved in practice.

2.4.3 Operational Lessons

The strongest conclusion we can draw from this survey of NAT operational methods is that a lack of a formal standard early in the development process enabled a market for consumer gateways which were defined by ambiguity in operation. Realizing the challenges imposed by these broad interpretations, the IETF attempted to clarify terminology and operational architectures with RFC 2663 in 1999 (later updated to precisely define behavioral requirements for UDP, TCP, and ICMP in RFCs 4787, 5382, and 5508 beginning in 2007) [169, 16, 62, 64]. This process of continual refinement continues with the most recent publication of NAT behavior requirements published in 2016 under RFC 7857 [139].

From a consumer perspective, these unclear device behaviors commonly challenge operation of services such as P2P sharing, online games, and voice-over-IP (VOIP) setup [16]. If a user had sufficient technical understanding, they could manually establish a rule within the gateway security policy to forward traffic originating from the internet to an internal device for the service in question. Depending on the type of NAT behavior employed, this could permanently open a "hole" into the customer's network, degrading any security afforded by NAT. In the worst case, options to fully expose a device exist within many gateway administrative menu's, often without warning to the consumer on the security implications [133]. To aid in managing this complexity, hole-punching methods commonly automate this configuration, removing the need for users to involve themselves in maintaining policy configurations.

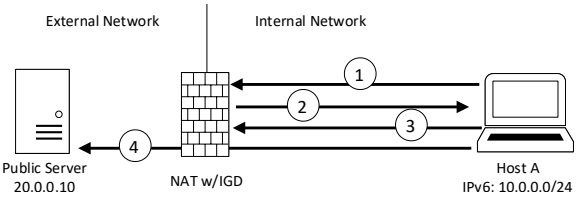

Figure 2.8: **UPnP w/ IGD.** 1) A new device executes a Simple Service Discovery Protocol(SSDP) request to identify supporting devices on the local network. Identified devices respond with a location (e.g., 192.168.1.1/service.xml) for the host to find defined services available. 2) Host request services listing through Service Control Point Definition(SCPD) to learn available actions to request. 3) Host requests an available action through Simple Object Access Protocol(SOAP), which instructs the IGD device to execute. In this case, it is a call to establish a port forwarding translation between the host and external interface. 4)The host informs a public server (or other external host) of how it may be reached for communication. The mapping is maintained until an explicit call to close the mapping occurs [18].

2.5 Overcoming NAT: Multiple NAT Traversal Methods for Multiple Behaviors

The default-deny behavior derived from NAT supported a simple and default security assurance to consumers. However, systems that required inbound connection establishment, such as VOIP, peer-to-peer, and others, needed a way to approximate the intended end-to-end design of communications. Mechanisms to "punch holes" through the NAT security boundary on behalf of the user provided this approximation. In many cases, these hole-punching methods rely on specific behaviors of NAT, resulting in an equally diverse and complicated set of solutions for the consumer to understand, deploy, and maintain.

In this section, we present a survey common hole-punching approaches, beginning with the most fundamental and commonly deployed mechanisms found in the majority of consumer gateways. For each sub method, we explain technical operation for the nearest canonical example and highlight related methods for brevity. Readers are encouraged to utilize associated references for a more detailed description of operational methods, as necessary.

2.5.1 Port Forwarding Methods

Port forwarding is a simple method for a user to statically map an external gateway port ($extIP{:}extPort$) to an internal host ($intIP{:}intPort$), enabling inbound communications across a NAT gateway. This mapping remains active until the user removes the configuration, potentially leaving a host exposed to unwanted communications if not properly maintained. To address these challenges of

user involvement and persistence, many automated mechanisms are widely deployed within consumer gateways, such as Universal Plug-and-play (UPnP) and port control protocol(PCP).

Universal Plug-and-play (UPnP)/ Internet Gateway Daemon(IGD) is a suite of discovery and coordination protocols which allow for seamless and automated gateway configuration, as shown in Figure 2.8. Here, a gateway daemon listens for local network participants to execute a configuration action. The permissive nature of who may initiate a configuration, combined with manufacturers enabling UPnP by default on many devices, has lead to many well-publicised security concerns. Notable examples include the "Unplug, Don't Play", "UPnProxy", and "CallStranger" UPnP attacks, which have exposed billions of consumer devices through improper implementation or flawed execution surrounding UPnP [118, 5, 20]. Despite these flaws, UPnP remains widely deployed, even at present. While consumers are advised to turn this feature off to limit security exposure, doing so requires direct involvement to disable - exactly what this protocol was meant to remove.

Port Control Protocol is the successor to NAT Port Mapping Protocol (NAT-PMP), a translation mechanism widely used by Apple systems. PCP works similar to UPnP, relying on server located on a NAT gateway to listen for and execute port configuration requests originating from the internal network [192]. Unlike UPnP, which is designed to enable management interfaces that allow for easy interaction by users, PCP is targeted to programmatic solutions that would typically be utilized by applications and computer programs.

2.5.2 Network Protocol Punching Methods

UDP hole-punching exploits NAT behavioral characteristics that allow inbound requests from any external host to be forwarded based on an active translation in the NAT forwarding table. As such, devices which use symmetric NAT behaviors cannot be used as traffic is restricted to both a single external host IP and Port. UDP punching is commonly found in peer-to-peer applications, VPN setup, and as a supporting method for tunneling mechanisms. This popularity likely stems from its broad success rate, with one study finding over 82% of consumer gateways presenting a successful traversal without requiring gateway configuration [54]. With UDP hole-punching, a publicly accessible server acts as a mediator to

coordinate connection establishment, as shown in Figure 2.9. If the connection is dropped, the hosts must re-establish communication by repeating the setup process. Keep-alive packets are commonly employed if a communication channel should remain active for an extended period of time [67].

ICMP and TCP hole-punching are distinct in that they are autonomous methods requiring no third party coordinator to trigger a path through a NAT gateway [141]. Due to the autonomous nature, setup requires strict coordination, prior knowledge of the endpoint gateway IP address, and predictable port selection in order to successfully execute [141, 76]. While shared knowledge of the destination IP address is easy to coordinate, the shared knowledge of which port will be generated is not [76]. Depending on vendor implementation of NAT, the selection of a port may be predictable using a simple known algorithm, direct mapping, or sequential selection [198]. If these port determination methods are not predictable, TCP hole-punching is unlikely to succeed. Second, operational differences with TCP connection handling may also prevent successful translation. For example, if a NAT gateway tracks an incoming TCP connection request destined for an active translation in the forwarding table, the gateway may drop the request completely or send an RST packet in response. his prevents the new connection from occurring, even though the same process may work with UDP [54]. This again shows that both the type of NAT forwarding, combined with unique device behaviors, play a large role in determining the best approach to establishing an active port forwarding in a gateway.

In a similar manner, ICMP hole-punching works by having an internal host send an ICMP Echo Request to an un-allocated remote address. In response, the NAT device will enable routing of replies, allowing an external connecting client to fake a "time-to-live: expired" message with their own address information. The NAT gateway sees this inbound client response as a match to the outgoing ICMP Echo Request, forwarding the packet to the internal host. This process allows protocols, such as TCP, to be tunneled over the UDP session, requiring no 3rd party setup or configuration to execute [121].

2.5.3 Tunneling Methods

For our classification, we define tunneling methods as any system or protocol that utilizes another to establish connectivity across a NAT device. These are loose definitions and aspects of other categorizations

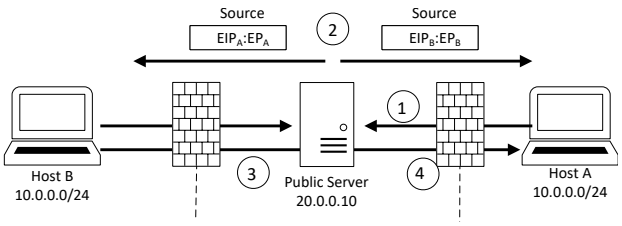

Figure 2.9: **UDP Hole-punching.** 1) Both host A and B establish communication with well-known public proxy about their intent to connect, resulting the each gateway establishing a port mapping back to the internal host, e.g.:($extPort_A : intIP_A$ 2) The public proxy server inspects both communication streams and forwards $extIP_A$:$extPort_A$ back to host B using its active connection with host B. It does likewise for host A. 3&4) Each host attempts to directly connect with the other using the active translation from each host's original request with the proxy server. This results in a new translation mapping in each gateway as follows: ($intIP_A$:$extPort_A$,$extIP_B$:$extPort_B$). Likewise, the same process occurs at B's NAT device, establishing two-way direct communication.

may play a significant role in establishing the following categorized methods.

LogMeIn/Hamachi uses a server-assisted NAT traversal technique similar to UDP hole-punching, but improves the methodology through a proprietary algorithm to increase success from 80% to greater than 95% [136]. In this server mediated method, each host initially establishes communication with an external moderator, as shown in Figure 2.10. The mediation server then instructs each host to conduct a NAT discovery probe, consisting of three separate UDP packets used to probe targets on the server (e.g., serverIP:port1, serverIP:port2, serverIP:port3). Information gained from these three probes is used by the server to better predict the port selection and type of firewall (stateless/stateful) operating on the local network. The information from these probes is then used to tailor the connection setup approach for behaviors of each gateway, thereby increasing the chance of success.

Teredo tunneling supports traversal of IPv6 clients located on private IPv4 networks. Conceptually, Teredo is very similar to Hamachi tunneling. First, a node, called a Teredo relay, acts as a gateway into an IPv6 network for which the tunnel for the IPv6 host will end. The mediation server in this scenario assists the client with establishing an IPv6 address, while the relay supports establishing an IPv4 UDP tunnel across the IPv4 network. This allows an IPv6 host to communicate across an IPv4 network and NAT device, even though they do not organically support such routing or traversal [78].

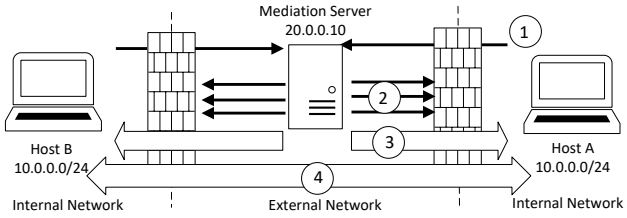

Figure 2.10: **Hamachi Tunneling.** 1) Both host A and B establish communication with well-known mediation server through local NAT gateways 2)The mediation server directs each host to execute a series of probes to learn about each gateway's forwarding characteristic and port assignment process. 3)The mediation server begins tunnel setup to each host and monitors for success. 4) If tunnel setup successful, the mediation server provides each endpoint with the other party's $extIP{:}extPort$ information and hands off tunnel so each host may communicate directly.

2.5.4 Client/Server and Relay/Proxy Methods

Proxying methods utilize an external, globally addressed, server to coordinate or assist endpoint hosts with NAT traversal. Interactive Connectivity Establishment (ICE) [146, 140, 106], Session Traversal Utilities for NAT (STUN), and Traversal Using Relays around NAT (TURN) are often grouped together as a single service due to ICE's use of TURN and/or STUN and the limited individual use of the latter protocols independently. These methods are commonly used to establish peer-to-peer (P2P) communications when both parties are located behind a NAT gateway. A common implementation of ICE is demonstrated in Figure 2.11. Here, two hosts (host A and host B) are ignorant of their own topology and how to best communicate with their remote peer. Each peer goes through a discovery process to identify potential candidate addresses and ports with which to establish a P2P session. These candidate addresses are then shared through a signalling channel, established via a publicly accessible proxy/signalling server, after which each peer begins a process of testing each remote peer address for connectivity [88].

While conceptually very similar to UDP hole-punching, we classify UDP separately to maintain alignment with the overarching focus on the core transport protocol. In practice, UDP hole-punching could also be classified as a relay method in which an intermediary is used to establish communication between two peers.

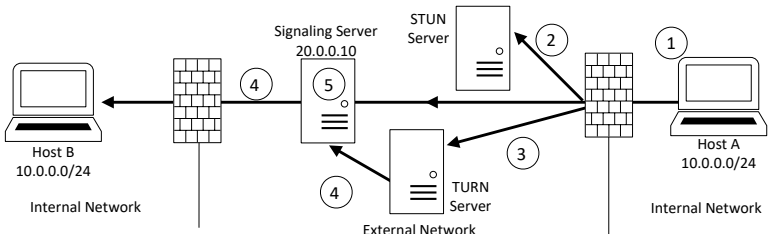

Figure 2.11: **ICE w/ TURN & STUN.** 1) Each host will go through a discovery process of learning all potential usable addresses in which to communicate with a remote peer. These candidate addresses will typically include connected interfaces (physical, virtual, or tunnel), 2) public facing gateway addresses discovered through STUN, or 3) if a TURN server is specified for relaying communications, it will receive an address and assign it as a candidate to use. 4) Each host then shares their list of potential candidate addresses with their remote peer via a signalling channel (via coordinating server/proxy), who will then test each available address until it finds a suitable candidate with which it can communicate. 5) Coordinating or signalling servers are commonly employed to coordinate ICE setup.

2.5.5 Increased Security Exposure for the Consumer

While many of these protocols increase the security exposure to a consumer network (which we show in Section 2.7), the IETF broadly supports this outcome. In assessing the NAT-PMP hole-punching method they state:

> The purpose of a NAT gateway should be to allow several hosts to share a single address, not to simultaneously impede those host's ability to communicate freely. Security is most properly provided by end-to-end cryptographic security, and/or by explicit firewall functionality, as appropriate. Blocking of certain connections should occur only as a result of explicit and intentional firewall policy, not as an accidental side effect of some other technology. This protocol goes some way to partially reverse that damage. However, since many users do have an expectation that their NAT gateways can function as a kind of firewall, any NAT gateway implementing this protocol SHOULD have an administrative mechanism to disable it, thereby restoring the pre-NAT-PMP behavior [26].

This position presents a number of troubling concerns with regard to consumer gateway security. First, the security exposures enabled by many of these automated abstractions are, to put it lightly, "a feature, not a bug". Secondly, this position both assumes and requires that users be active participants in precisely managing their own security policies. This position counters the efforts by manufacturers to simplify and abstract security *away* from the user [178, 142]. Third, the concluding position hesitantly recommends an administrative mechanism to restore the assumed benefits of NAT through an active interest and involvement by a user rather than as default guarantee.

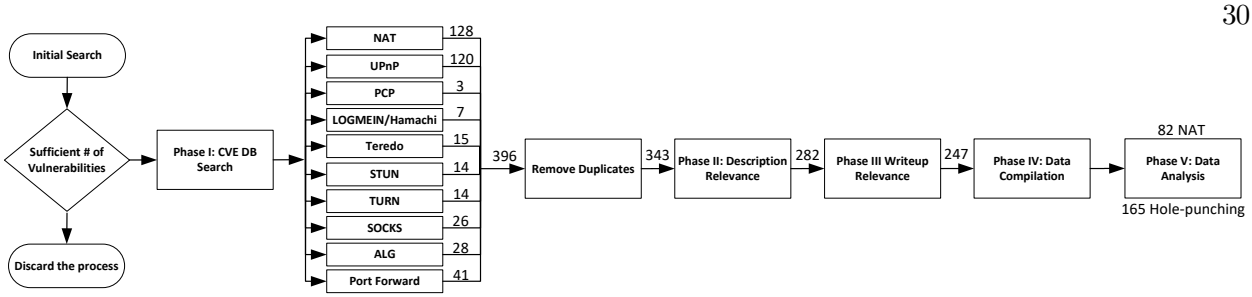

Figure 2.12: **CVE Review Process Map.** Relevant hole-punching protocol technologies were searched and sorted to assess overall takeaways. In total, 82 CVEs related to NAT and 165 related to hole-punching methods were used in our assessment.

While the IETF has acknowledged the unique challenges of balancing consumer network security with broader Internet architectural goals, positions in favor of end-to-end connectivity and hesitation to both define and implement controls counter to this design remain [138]. Recent efforts to assess IPv6 operation within consumer gateways present many of the same challenges and pitfalls [133]. The result is an ambiguous operational environment of technologies within consumer gateways that have no clear security or operational guarantee.

2.6 Taxonomy of NAT and Hole-Punching Method Security Flaws

In the preceding survey of operational methods for NAT and associated hole-punching methods we present a theme demonstrating how ambiguity or an absence of a defined standard has lead to a diverse and challenging operational environment, not only for developers, but consumers as well. In the following sections we narrow our focus and assess how these mechanisms intended to ease consumer management have traditionally degraded the overall security within a consumer gateway. In order to align our efforts with previous works, we adopt the literature review and organizational methods utilized in [91, 112].

For our review, we rely on the NIST National Vulnerability Database (NVD) [129], the MITRE Common Vulnerabilities and Exposures (CVE) [114], and the US Computer Emergency Readiness Team (US-CERT) Vulnerability Notes databases. We performed a vulnerability search and selection process over five steps, demonstrated in Figure 2.12.

We began by first conducting a cursory search to determine if sufficient vulnerabilities exist to

conduct the survey of NAT and hole-punching security failures. Then, using broad search terms, we gathered over 300 documented vulnerabilities representing exposures within both consumer gateways and many commercial implementations[2] Duplicate entries, resulting from reliance on two database, were removed and a CVE vulnerability description review was undertaken to validate relevance of each result. When relevance could not be obtained through the vulnerability description alone, review of the supporting documents were conducted to determine final selection or rejection. In total, we identified 82 vulnerabilities directly related to NAT and 165 vulnerabilities related to hole-punching methods for our analysis.

2.6.1 Hole-Punching and NAT Security Taxonomy

The CVE 2.0 categorizations provided by the National Vulnerability Database (NVD) present a common classification reference, framing a vulnerability in terms of complexity, impact effect, and severity. While these categorizations serve to assist with gauging the relative impact of a vulnerability, our goal is to survey the security failures as a whole in order to understand the breadth of exposures present and the mechanisms leading to their failure. To do this we conducted a three step search and review process focused on discovery and categorization as follows:

(1) For step 1, we reviewed each CVE, documenting unique characteristics that aided in defining an attack. From this review process, we arrived at the following classification categories: vulnerability relationship to assessed protocol, network source of attack, primary security flaw or weakness, primary effect of exploiting the weakness, resulting exposure, and overall impact to security.

(2) For step 2, we relied on the identification of traits from step 1 but further grouped each assessed vulnerability into sub-categorizations, arriving at the taxonomies presented in Figure 2.13.

(3) For step 3, we conducted an additional review to ensure all identified vulnerabilities were accounted for and aligned to a category within our framework, thus validating and ensuring completeness of our process.

[2] We expand our search beyond just consumer gateway security failures in order to broadly capture the failures within these mechanisms. While a commercial device security failure does not directly represent an exposure within a consumer network, their failures commonly represent exposures which *could* occur within a consumer gateway and the mechanism by which a failure may occur. Therefore, where appropriate, we maintain these exposures in our analysis to provide broader insight into the weaknesses presented by these mechanisms.

2.6.2 Taxonomy Category Classifications

In the following section, we present our taxonomy classification categories, along with related statistics resulting from our analysis, which we use to draw security takeaways.

2.6.2.1 Classification Based on Target Relation (T)

Target relation classifications define the main relationship of the assessed protocol to an attack outcome. Within this category we identify and define three distinct classifications: failures related to protocol implementation, failures aided by protocol use, and for NAT, a third category of tangential failures occurring in other systems resulting from the use of NAT.

Protocol Implementation (T-1) flaws are underlying weaknesses in implementation that are directly associated with either NAT or hole-punching methods. In our analysis, 60% of identified security vulnerabilities for hole-punching methods fall within this category, while 40% of NAT vulnerabilities are directly related to implementation within a system.

Protocol Aided (T-2) flaws are second order security exposures which occur elsewhere within a device or network resulting from the use or operation of the assessed methods. When categorizing protocol aided security events, we focus on the final security exposure resulting from identified security weakness. For example, one deployment of UPnP within the Linksys WRT54G gateway allowed remote attacker's to arbitrarily forward ports on the system due to no implementation of an origination validation process for a "addPortmapping" request [35]. Protocol aided flaws represent the remaining 40% of identified flaws surrounding hole-punching methods while only representing 7.3% of NAT flaws.

Tangential (T-3) are flaws in other systems that occur through incorrect handling or processing resulting from the assessed method. This identification only occurred within the NAT environment, often exposing a system through incorrect use of public/private addressing or improper access control for NAT'd systems. Tangential flaws comprised 54% of all security flaws related to NAT.

2.6.2.2 Classification Based on Network Location (NL)

Network location categorizes an attack based on the vantage from where a successful exploitation can occur. Classification of network location is directly obtained from each CVE and are used to conduct overall trend analysis in Section 2.7.

Adjacent (NL-1) attacks originate from within the network boundary in either the same subnet, collision, or broadcast domain. A common example scenario would be a WiFi broadcast domain such as a coffee shop or other shared access environment. Attacks requiring network adjacency make up 9.7% of hole-punching and 3.7% of NAT vantages.

Remote (NL-2) attacks originate from an external network, typically one or more network hops away. Remote attacks require use of the OSI Network layer for execution. Across both assessed methods, remote vantages represent the most common exploitable vantage, representing 82.4% and 87.8% of occurrences for hole-punching and NAT, respectively. We suspect these classifications within the NVD over-rely on the "remote" classification of attacks due to term usage ambiguity, which we discuss further in Section 2.7.3.

Local/System (NL-3) attacks require direct access to the target device in order to successfully execute. These type of attacks commonly exhibit a flaw in code or resource management that cannot be exploited through remote interaction. Together, local/system access represent 7.8% and 8.5% of security flaws for hole-punching and NAT respectively.

2.6.2.3 Classification Based on Security Weakness (SW)

These classifications are defined by the primary failure of a system or protocol that leads to an exposure. Within the National Vulnerability Database, vulnerabilities are assigned a weakness enumeration value corresponding to one of hundreds of possible weaknesses. In cases where multiple overlapping definitions occur, we have consolidated them into a single category to focus on the broader security concern (eg. CWE-119 "Improper Restriction of Operations withing the Bounds of a Memory Buffer", CWE-120 "Buffer Copy Without Checking Size of Input", and CWE-121 "Stack-based Buffer Overflow"

are classified as a single "Improper Buffer Restriction".)

Additionally, security flaws may build upon one another, resulting in a sequence of exposures that lead to an eventual compromise. As an example, weak input validation may lead to a buffer overflow condition which results in the ability to perform a code execution attack. For our categorization process, we focus on the initial flaw as the primary security weakness for our categorization as it is the root vector.

Improper Restriction of Buffer (SW-1) occurs when an operation extends beyond its assigned bounds within memory. Buffer overflows are the most common type of security weakness found in hole-punching methods, resulting in nearly one quarter of all exploitation effects. In contrast to hole-punching mechanisms, weaknesses within NAT resulting in buffer overflows are the least common security weakness, occurring in less than 3% of reported security flaws in our assessment.

Input Validation (SW-2) failures improperly check user inputs against expected values or length. While improper input validation is a common vector for buffer overflows, we differentiate this categorization when the input validation failure is the primary avenue or method to initiate an exposure resulting from a user-provided input. Input validation security weaknesses are common to both assessed methods, occurring in 17.7% of hole-punching and 17.1% of NAT weaknesses.

Permissions, Privileges, and Access Control (SW-3) are a broad categorization of many security weaknesses that fail to restrict access or device interaction to an authorized scope, resulting in exposure of a device, controls, or data. This categorization has the second highest rate of occurrence within hole-punching methods, accounting for 20.7% of assessed weaknesses. This occurrence drops significantly within NAT, accounting for only 4.9% of assessed security weaknesses.

Resource Management (SW-4) weaknesses result in uncontrolled utilization or improper bounding of a system resource. For example, NAT implementation within versions of the Cisco IOS resulted in memory leaks via malformed SIP packets attempting to traverse a gateway [36]. Resource management flaws occur in 22% of assessed vulnerabilities for NAT and 10.4% of hole-punching weaknesses.

Improper Credential Authorization, Bypass, Protection (SW-5) flaws are the result of an attacker obtaining elevated access to a system through improper presentation and acceptance of credentials by a system, or by bypassing authorization mechanisms which restrict user access. Authentication flaws

occur in 12.2% of hole-punching and 4.9% of NAT assessed security weaknesses.

System Configuration (SW-6) weaknesses are those in which the default configuration of a device fails to present a secure operational baseline. As an example, this categorization could result from configurations where services intended for use on an internal network are improperly configured to operate on the untrusted side of the network, which occurred in the commonly referenced "Unplug, Don't Play" Rapid 7 assessment of consumer gateway security [118]. System configuration flaws occurred in 8.5% of hole-punching and 2.4% of NAT assessed weaknesses.

Coding Error (SW-7) encompasses the many potential methods in which a program may fail where a more specific categorization is not present, such as with an off-by-one calculation error. Coding errors present a small, but unique, subset of weakness classification, representing 4.8% and 7.5% of assessed weaknesses within hole-punching and NAT.

Clear-text Transmission of Sensitive Information (SW-8) presents just one example (0.6%) within the hole-punching categorization. Here, a device presented administrative credentials to any adjacent user performing a UPnP "X_getAccess" SOAP request to the Internet Gateway Device(IGD) [42].

Improper Resource Validation/Handling (SW-9) occurs within systems that fail to properly check or account for varying responses to processing inputs. For example, the Windows implementation of NAT in Server 2012 did not properly validate memory addresses when processing ICMP packets, resulting in a denial of service condition [37]. This type of weakness is commonly found within the NAT processing environment, where packet processing implementations fail to account for address translation, commonly resulting in unintended exposure of devices and networked systems. This flaw is the most common security weaknesses within NAT, representing 23.2% of assessed weaknesses. There were no resulting weaknesses identified for this category within the hole-punching classification.

2.6.2.4 Classification Based on Primary Effect (PE)

Primary effects result from the exploitation of a system weakness. They represent the final goal an attacker would seek to achieve.

Denial of Service (PE-1) occurs when a device is no longer able to service legitimate requests.

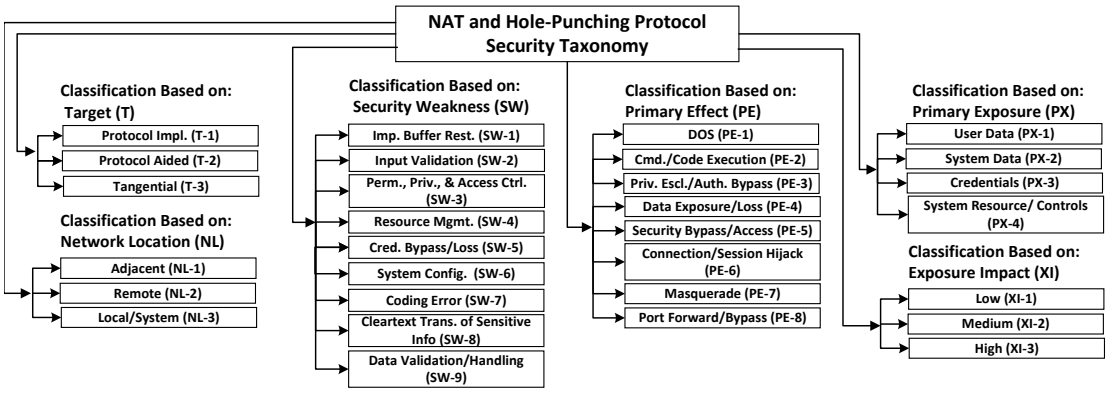

Figure 2.13: **Final Taxonomy of Hole-Punching and NAT Security Failures.** The taxonomy is based off systematic review of 300 CVE documents to obtain source classifications of security flaws based on target, network vector, security weakness, attack effect, and resulting exposure. A single taxonomy is presented to represent an overall assessment of gateway security flaws and due to the significant overlap found in conducting the taxonomies separately.

Common methods include system crashes due to buffer overflows, resource exhaustion, or configuration changes resulting in a service outage. Denial of service is the most common outcome for both hole-punching and NAT effects representing 29.1% and 64.6% of assessed effects.

Code/ Command Execution (PE-2) is one of the most critical vulnerabilities as it allows an attacker to change the behavior of a system. Devices, such as VeraEdge, have demonstrated attacks in which the UPnP service accepts un-sanitized URLs, enabling code execution via a buffer overflow. A number of buffer overflow flaws in UPnP alone allow attackers to execute code on a local device [118]. This is the second most common effect within hole-punching vulnerabilities, representing 17.6% of vulnerability outcomes. Only 6.1% of NAT vulnerabilities experience this effect.

Authentication Bypass/Privilege Escalation (PE-3) are effects which provide an attacker some level of access to a targeted system. These effects are commonly found within the hole-punching category as many of the methods provide avenues for an attacker to interact with and exploit the targeted device by bypassing authentication controls. 13.9% of hole-punching effects provide some level of privilege escalation or bypass. In contrast, only one instance of NAT allowed for an attacker to obtain elevated privileges based on an application improperly relying on a gateway address for device identification, resulting in all NAT'd users being provided administrator access [39].

Data Loss/ System Information Exposure (PE-4) is a broad categorization of exposures

resulting in an attacker accessing or viewing information reserved for a privileged or restricted scope. The attacker is not able to execute any further direct attack beyond the viewing of exposed privileged information, though the information may enable further efforts such as direct targeting of a device. This effect is the second most common outcome for NAT vulnerabilities, representing 14.6% of assessed exposures. For hole-punching, only 7.9% of vulnerabilities exhibited this outcome.

Security Bypass/System Access (PE-5) Any method in which the primary effect presents access to the system in which an attacker may execute further action are presented under this category. This categorization extends beyond the authentication/privilege bypass methods previously categorized by focusing on system level flaws that enable access to a targeted device. Vulnerabilities exhibiting this effect occur in 11% of NAT and 4.9% of hole-punching classifications.

Connection/Session Hijack (PE-6) occurs when an attacker is able to take over control of an active connection/session. For NAT, two occurrences of a session hijack occur. In the first case, a sip registration service failed to properly require registration when NAT was enabled, allowing a remote user to take over any active session [40]. In the second case, a Netgear DIR-615 router identified users by their gateway IP for remote access, allowing an attacker to sniff the gateway public IP and take over a session without being prompted for credentials [41] For hole-punching, incorrect implementation of the TURN/STUN protocol within WeMo devices allowed an attacker to hijack connections to any other connected WeMo device [38].

Masquerade (PE-7) differs from a connection hijack in that the attacker is able to establish their own connection under another user or session. This effect again presents itself rarely, representing just a single occurrence across both NAT and hole-punching effects.

Port Forward (PE-8) is unique only to hole-punching methods. Port forwarding is a desirable effect to an attacker as it provides a path for inbound traffic to traverse a perimeter security implementation, such as a firewall or NAT gateway. Port forwarding represents 14.6% of assessed security effects within the hole-punching category.

2.6.2.5 Classification Based on Primary Exposure (PX)

Primary exposures define the primary type of data or access revealed by an attack. The CVE classification methodology relies on the familiar CIA triad of confidentiality, integrity, and availability when categorizing an exposure, with sub categorizations of none, partial, and complete (None, Low, and High for CVE 3.x). While this methodology provides for a quick assessment of impact across the core tenets of information security, it does little to communicate what exactly is being exposed. Therefore, we expand on this classification, identifying from our dataset four categorizations of exposure that identify what an attacker may ultimately gain.

User Data (PX-1) consists of all data generated by a user and may include items such as payload data in IP communications, metadata such as use statistics, or identification of devices within an environment. One example of this type of data loss would be the public exposure of IP cameras which allowed a remote attacker to eavesdrop via publicly exposed STUN ports [109].

System Data (PX-2) exposure consists of device information such as type, configuration, or protocol communication traffic which could be used to fingerprint or determine exposure to known vulnerabilities. This information typically provides information that enables follow-on targeting of system components.

Credentials (PX-3) are any event where the primary effect results in the attainment of system or user credentials. Methods to bypass credentials are not classified here as they would provide direct access to system resources or control.

System Controls/Resources (PX-4) are those in which any unauthorized user is presented with access to a device or protocol control or resource. Attacks in which malicious users are afforded this type of exposure typically result in changes to operational state or configuration in ways that are beneficial to the attacker. This may include methods to further goals, such as with code injection, or as simply an end means, such as corruption of resources.

2.6.2.6 Classification Based on Exposure Impact (XI)

Exposure impact communicates, in broad terms, the potential impact to a user, device, or network resulting from an attacker successfully exploiting a weakness. There are two classification methodologies present in the NVD, the CVE 2.x methodology and the CVE 3.x methodology. The 2.x methodology classifies an impact as either a Low, Medium, or High threat while the 3.x expands this classification to include None and Critical categories. The 3.x methodology was first introduced in 2016, limiting applicability across all of our assessed vulnerabilities. However, the NVD continues to provide 2.x scoring along with the newer 3.x deployment, allowing for direct comparison of vulnerabilities and trends. For our impact classification, we rely on the 2.x categorization of impacts provided by the NVD, to allow for direct comparison across all vulnerabilities.

Low (IX-1) represents a CVE impact scoring of 3.9 or less. When reviewing NAT and hole-punching methods, a total of four NAT and eleven hole-punching impact scores fell in this categorization, representing 4.8% and 6.6% of the total assessed vulnerabilities.

Medium (IX-2) represents an impact score ranging from 4.0 to 6.9. A total of 31 NAT and 69 hole-punching methods received a Medium score, representing 37.8% and 41.8% of the total vulnerabilities assessed.

High (IX-3) represent the greatest impact categorization covering scores between 7.0 and 10.0. A total of 47(57.3%) and 85(51.5%) examples fall within this categorization for NAT and hole-punching, respectively, representing both the largest share of events and greatest threat to a user or network.

2.6.3 An Increased Impact to the Consumer

Revealed by these taxonomies are the significant exposures occurring within consumer gateways via mechanisms intended to ease access control management away from the user. Further, the security value of a default-deny perimeter policy provides little value to a consumer when mechanisms to circumvent commonly introduce far greater risk. These additional exposures are not relegated to minor considerations. Over 50% of the assessed vulnerabilities carried a "High" risk rating. When combined with the number

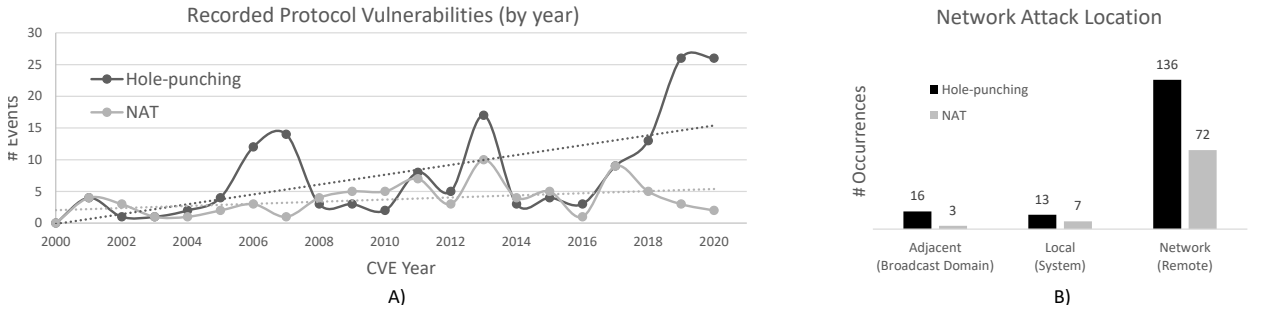

Figure 2.14: **CVE Distribution by Year and Primary Network Avenue of Attack.** A) Documented vulnerabilities for hole-punching methods have increased over time, in line with broader security trends overall. Contrary to this increase, NAT has remained relatively unchanged, averaging 1.8 CVE's per year. Assessment is based on directly related CVE's only. Tangential or protocol aided attacks are not included in this distribution. B) CVE scoring shows a disproportionate number of vulnerabilities remotely exploitable from external network vantages.

of flaws surveyed, over three-hundred, this begs a question on whether the inclusion of many of these protocols meant to aid a user actually provide any value at all.

These exposures should also highlight the need for revisiting the default enablement of many of these mechanisms and what exactly should be included in a baseline security definition for home gateways. While we cannot tell how many users rely on these aids, we believe that an opt-in approach is the necessary and correct answer. In practically all cases, the solution to address the security flaw would require a firmware or software update. Studies assessing the frequency and completeness of these updates show very little effort by gateway manufacturers to address and if so do so in a timely manner [68]. However, there is little incentive for manufacturer's to improve this present state as market factors commonly outweigh the effort needed to establish stronger security baselines [149]

2.7 Analysis of NAT and Hole-punching Common Vulnerabilities and Exposures

In the prior sections we introduce both the breadth of available access control mechanisms and a taxonomy of security failures present within these mechanisms. We continue this analysis by investigating both the historical trends over the life-cycle of these access control mechanisms and the statistics surrounding the security exposures ultimately introduced into the consumer network.

2.7.1 Vulnerabilities Over Time

Ideally, a system or software package will enter the market in a thoroughly tested and reviewed state. However, differences in implementation, proprietary development, or trailing standards allow opportunities where security is likely to fail. Further, incentives for "first mover" or "first to market" encourage inclusion of systems or components that may not yet be standardized or fully tested [149, 143]. In an ideal, mature process, these security shortcomings would generate a cycle of patching that builds toward a secure steady state. What we find in our analysis is a dichotomy between NAT and hole-punching protocol vulnerabilities over time. Figure 2.14 shows that over the life-cycle of hole-punching methods there has been a steady rise in discovered vulnerabilities. This growth is in line with general trends in CVE reporting overall [98]. In contrast, NAT demonstrates a slower growth in documented exposures, averaging roughly 4.1 vulnerabilities per year. We posit four reasons for this disparity:

(1) NAT is integrated into the Linux kernel via the Netfilter package library. Nearly 90% of home gateways use the Linux kernel for implementing core OS functionality [189]. This commonality allows for a single package maintenance across nearly all gateways. In contrast, there is a wide variance in packages used by manufacturers to implement hole-punching methods. For UPnP, there are over 1,500 unique implementations available on GitHub, though only ten of these represent 90% of deployed instances, excluding versioning [118].

(2) The codebase for core NAT functionality in Netfilter is roughly one thousand lines of code(LoC). In contrast, the complete package for MiniUPnP is over forty-five thousand LoC [174]. With an average of fifteen to fifty bugs per one-thousand LoC, the potential for mistakes in hole-punching packages increases significantly [110].

(3) Devices are not being readily patched, allowing for discovery of additional vulnerabilities across package versioning. To quantify this later point, nearly 25% of MiniUPnPd deployments worldwide still use version 1.0 despite over twelve major package releases addressing significant vulnerability concerns [180].

(4) Updating separate software packages can be costly from a development perspective, as changes

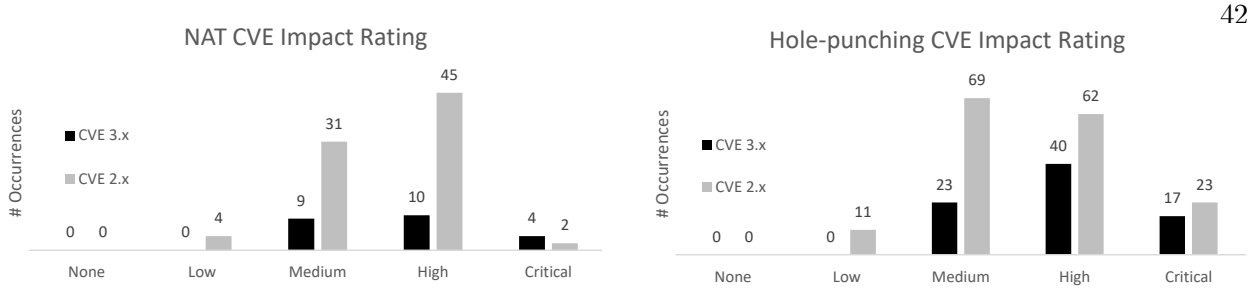

Figure 2.15: **CVE Scoring Distribution - NAT and Hole-punching Protocols.** Scores for CVE 2.x and 3.x are displayed. Scoring for 2.x represents *all* vulnerabilities while the 3.x reflects only the CVEs which have been scored under the newer metric. The CVE 2.x scoring is aligned to the 3.x categories based on score only. For NAT, the average score under the 3.x metric is a 7.56 or "high" classification, while average score under 2.x is a 6.56 or "medium" classification. For hole-punching methods the average under the 3.x metric is a 7.71 or "high" classification, while average score under 2.x is a 6.40 or "medium" classification.

may introduce 2nd and 3rd order efforts to ensure a new software package is compatible with the overall system. Therefore, there is little incentive for a manufacturer to actively maintain these packages.

2.7.2 CVE Severity Distribution

The CVE severity scores reflect the severity of each documented exploit on a scale of zero to ten, with ten being the most severe. Within this distribution scale are sub-categories of None (score of zero), Low (0.1-3.9), Medium (4.0-6.9), High (7.0-8.9), and Critical (9.0-10.0). Two methods of scoring were used for our analysis: the older CVE 2.x covers *all* of the documented attacks while the newer CVE 3.x, introduced in 2016, covers approximately half of the documented exploits. The first point of interest is that the scoring between 2.x and 3.x skews severity classification higher under the 3.x model. This is in-line with general comparisons between 2.x and 3.x scoring overall [152]. When like metrics are compared for vulnerabilities that have scores for both methods, the average hole-punching vulnerability for 3.x scoring is 7.71 compared to the average 2.x scoring of 6.40. This is a critical point to make clear as the lower represents a medium threat, while the higher represents a high threat for the same average vulnerability. Second, and the larger point of concern, is that the average vulnerability surrounding hole-punching methods represents a high threat to consumer security overall.

Similarly, NAT exhibits the same high severity classification. The average 3.x scoring results in

a 7.56, or high rating, while the CVE 2.x scoring for the same vulnerability average results in a 6.56, or medium classification. Of note are a disproportionate quantity of vulnerabilities with only 2.x scores due to the majority of NAT vulnerabilities occurring prior to the shift to the newer scoring standard. The distribution of each is shown in Figure 2.15.

Despite these threats to security within the home, research has shown that upwards of 60% of users run outdated firmware within their home gateways [181], representing a significant exposure to security vulnerabilities within the home network. Recent efforts by manufacturers to address this gap now include automatic updates to ease consumer burden [10, 61]. However, it is unclear whether these automatic updates actively maintain all component software packages or if they fall into a similar trap of patching only significant exposures. In either case, the frequency of updates offered by most manufacturers can span months to years, allowing significant time for gateway exploitation [68].

2.7.3 Access Vector Analysis

Network access vector defines the type of presence required to execute an attack. The CVE scoring system uses three classification categories to define access: (N) Network attacks are those which are realized at layer 3 or above of the OSI network stack and an attacker does not require local network access. These could commonly be considered remote attacks. (A) Adjacent vantages are those in which the victim and attacker are on a shared network segment, such as a shared broadcast or collision domain. (L) Local or system level access requires an attacker to have access to the machine at hand, either through physical access or a local account. Of greatest concern would be a remote attacker who is able to exploit a vulnerability on a target.

Across the CVE scoring system nearly 83% of all exploits were documented as network exploits, shown in Figure 2.14B. We believe this to be the result of unclear definition in the original CVE 2.0 standard, likely resulting in many exploits being improperly classified. For example, the Macintosh iChat UPnP buffer overflow is listed as a network access vector even though the description highlights a need for a local or adjacent access [43]. Similar examples exist throughout the network vector classifications.

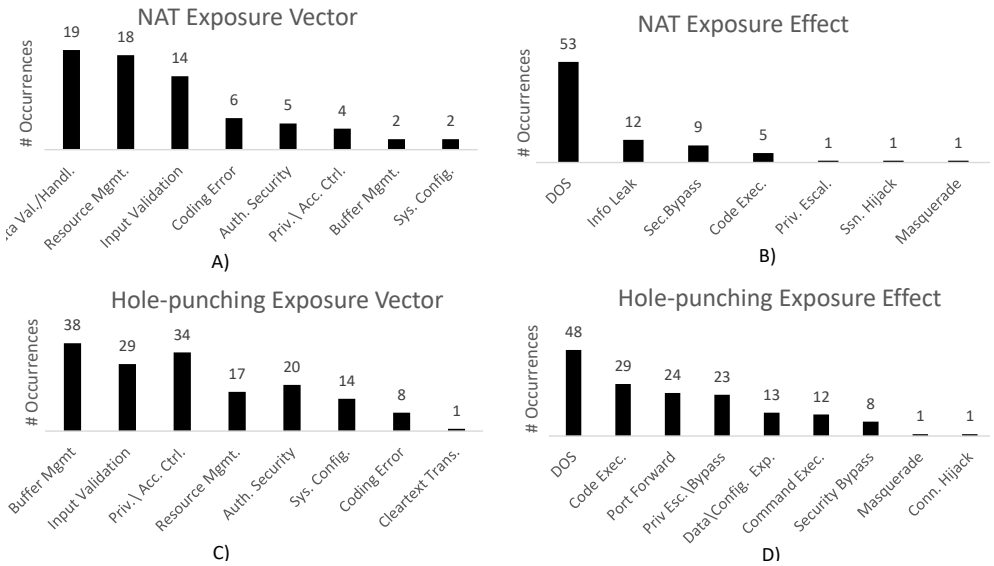

Figure 2.16: **Enumerated Weakness Vectors and Effects.** Enumerated occurrences of weaknesses vectors and resulting effects within NAT and hole-punching methods are presented. Note: hole-punching enumeration does not include six undisclosed vulnerabilities that could not be categorized.

2.7.4 Enumerated Weakness Analysis

Weaknesses are the exploitable flaw within the protocol implementation, which allow an attacker to achieve an effect. Across the breadth of NAT and hole-punching implementations we see significant exposure to the consumer, presented in Figures 2.16. Here we highlight both the critical weakness occurrence and the resulting exposure to the consumer. In one critical implementation failure, the UPnP daemon was exposed and operating on the external interface of many home gateways allowing remote attackers to create forwarding rules which allowed access to internal networks [118]. Globally, nearly 450,000 devices still maintain this exposure eight years later [160].

More concerning are the effects that an exploitable vulnerability may reveal. Nearly one-third of all exploits in hole-punching methods result in a denial-of-service condition. This ratio nearly doubles under NAT. While exposure to this type of effect would inconvenience the consumer with lost connectivity or productivity, violating the core principle of availability, a user's exposure is likely limited. In contrast, code execution, privilege escalation, and port forwarding vulnerabilities do significantly expose a user and are common within the realm of potential effects related to hole-punching methods. Together they comprise

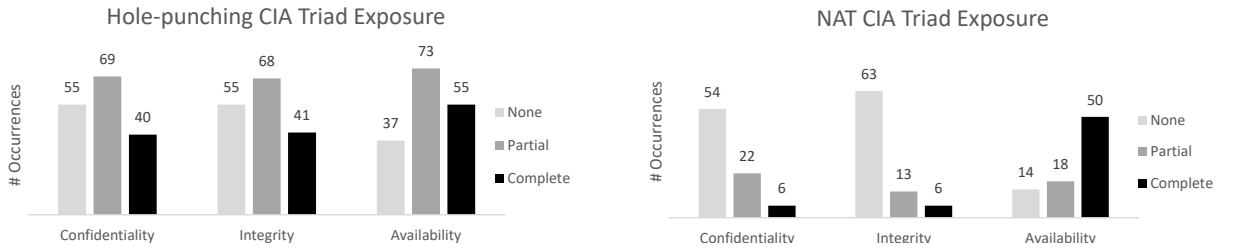

Figure 2.17: **CIA Triad Exposure - Hole-Punching Methods (Left) and NAT (Right).** Exposures to confidentiality, integrity and availability of systems due to exposures within or created by either NAT or hole-punching methods are presented. Exposures are ranked on a partial, complete, and none categorization representing the degree of exposure from an attack.

nearly one-half of attack effects, demonstrating a significant level of exposure to the consumer overall.

2.7.5 CIA Triad Exposure

Within the CIA Triad, user exposures are significantly higher within hole-punching methods, as demonstrated in Figure 2.17A; a testament to the increased risk presented by mechanisms meant to ease consumer involvement. Many of these automation aids are pre-enabled by manufacturers, resulting in immediate exposure to the consumer [101, 118].

Contrary to these extreme exposures within hole-punching methods are the exposures related to NAT, shown in Figure 2.17. In the majority of NAT vulnerabilities, weaknesses and attack effects result in little to no exposure to a user beyond loss of availability. This limited exposure begs the question: if the inclusion of hole-punching mechanisms result in significant security exposures, should they even be included within a home gateway in the first place? A consumer could enact the same functionality through carefully managed firewall policy that is aided by clear configuration options without the increased exposure presented by operating many of these management abstractions.

2.8 Discussion

Historically, home router security has never been good [131]. Systems have been found to implement software packages over ten years old and in many cases never fix exposures at all [68]. Incentives to force better security within these gateways remains notably absent. At best, the single strongest incentive for

manufacturers to provide security improvements is via future hardware upgrade sales. Notably, this also *disincentivizes* current hardware support as personnel resources are commonly aligned to future product developments. Solutions, such as device labelling standards or contractual support guarantees, have been proposed as possible mechanisms to incentivize better security through improved consumer awareness [125, 119]. However, given the lack of consumer involvement at present, these are likely to offer little value.

The single strongest incentive to force change is through regulation. Demonstration of regulation in the SOHO environment as a necessary means has already occurred, as highlighted by California enacting state law SB-327 in 2017. This law was enacted to combat the insecure default configurations of gateways, requiring manufacturers to no longer use default credentials on devices sold within the state [81]. The law further requires that mechanisms must be present to force a user to create a unique password on initial device setup. Unfortunately, this law does not go far enough as there is no stipulation for non-compliance. In reviewing present implementations for gateways built after 2022 we found that less than 10% of manufacturers actually complied with this law, showing that more regulation is likely necessary to force broader security changes within the consumer gateway industry. The present incentives have done little to move the bar forward for security over the last thirty years.

2.9 Related Work

Given the long life of NAT, a number of surveys and taxonomies have been conducted by researchers attempting to catalogue everything ranging from the operational characteristics to traversal techniques and security flaws. Many of the surveys surrounding NAT core behaviors and operational architectures are best documented in early efforts by the IETF in [169, 75, 16, 64, 62, 139]. These are supported by academic efforts which further identify and define device specific behaviors related to unique vendor implementations [120, 71, 63] and the IETF's response to help standardize these behaviors [139]. Considerations for specific architectures include ISPs [163, 144] and home networks [120, 71]. Similar efforts to provide holistic assessments on NAT operation related to specific protocols, such as SIP or IPSec, are commonly present in many of the early classification surveys (pre-2010) [25, 34]. However, we could not find any efforts to revisit these early classifications for correctness or in light of new technology implementations, such as IPv6.

Supplementing these behavioral classifications are works that consider NAT traversal techniques. Efforts focused on Peer-to-peer [54] and VoIP[79, 172] provide for a broad survey of approaches. For our work, we build on these early surveys by including and classifying the primary traversal mechanisms used within consumer gateways and further provide updated references and techniques within our classification.

Efforts which consider the role of NAT as a security mechanism are first discussed in [166] with many subsequent works [175, 83, 12] ultimately questioning this position. These are further supported by retrospectives on the role of NAT in particular, with many identifying the missed assumptions and poor standardization early as sources of continued challenges [197].

In considering the security of the gateway itself, a number of efforts have provided demonstration on specific topics related to the home gateway. For layer 2, security considerations and techniques for Ethernet [90] and wireless [186, 191] both provide recent a detailed analysis of access control mechanisms and attacks surrounding WiFi. Authors of [128] provide a more holistic assessment of gateway security challenges.

While this work maintains a narrow focus on consumer gateway access control and the resulting security challenges, consideration of additional controls within the customer premise are equally important to providing holistic security measures. Here, a large number of current surveys papers for IoT and Smart Home devices provide both a comprehensive and detailed review that would make their inclusion in this work of limited value [73, 123].

In reviewing these prior works, we found no singular effort provided for a broad assessment or holistic categorization of behaviors, mechanisms, or security concerns. Of further note is the absence of assessments that combine these topics the context of the most common usage of NAT and hole punching mechanisms: the consumer gateway. In light of these absences, we further provide a "long view" look at the security impacts presented by these mechanisms over time and relate these considerations to IPv6 deployment.

2.10 Conclusion

After thirty years of home networking, consumer gateways still rely on the simplistic model of a network perimeter first established by NAT. While arguably not a strong security solution, the default deny architecture undeniably provided a host of privacy and security benefits for non-technical users

via a near-universal operational baseline. Address scarcity served as an incentive for manufacturer's to deploy address translation, ensuring broad adoption and a common design architecture.

Unfortunately, the addition of mechanisms meant to further abstract consumer involvement have demonstrated a clear detrimental effect that runs counter to the purpose these abstraction mechanisms were meant to provide. The end result is an increased exposure of home networks that consumers are ill-equipped to address. Worse, manufacturers have little incentive to address these exposures arguing it is the consumers requirement to enact stronger security, a position in direct opposition to their efforts meant to remove the user's involvement.

Looking forward, we see the underpinnings of many of the same mistakes occurring, particularly with IPv6 deployment (which we assess in detail in Chapter 3). At present, open-ended requirements for IPv6 operation within consumer gateways enable a security environment defined by ambiguity, similar to the early days of address translation deployment. The potential for manufacturers to deploy two very different security models under IPv6 presents a challenge not only to the home user, who may not have the technical skills necessary to appropriately address, but also the manufacturers. This is particularly true for IoT and smart home devices, where many lack the necessary security mechanisms to perform precise access control themselves due to low cost designs and reliance on trusted operational environments within the home. Manufacturers, in particular, have a much larger responsibility, as many security flaws can be directly attributed to poor implementation or maintenance practices which are incentivized by market dynamics and product cycles. It appears, at present, that there is little incentive here forcing manufacturers to address these shortcomings with newly released systems still relying on long outdated security software [134], which they are slow to patch, if ever [23].

Absent stronger standards or policies, or a more involved consumer, there is little here to incentivize the change needed for stronger security guarantees.

Chapter 3

How are Manufacturers Incentivized to Apply Security Moving Forward? A look at IPv6 Filtering

Customer edge routers are the primary means by which most consumers connect to the Internet. As these consumer networks migrate from IPv4 to IPv6, the use of address scarcity to incentivize a common operational baseline and default-deny security posture is no longer present. On one hand, manufacturers could choose to follow the end-to-end principle [2] for networks, removing security from the gateway all-together. This approach may in fact align with business incentives to reduce cost and complexity for users, pushing security mechanisms and controls to the end devices. Alternatively, deployment of a stateful firewall could provide the same assured baseline consumers have grown accustomed to with NAT, assuming a set of default policies are appropriately implemented. Whether or not business incentives to reduces costs and complexity for users outweighs the need for familiar security baselines and user controls is unclear.

Assuming manufacturers continue to view the customer edge as a critical security demarcation by choosing to implement stateful firewalls for IPv6, the IETF provides little guidance for directing the use, configuration or baseline requirements of such an approach [161, 28]. While the use of standards could serve as a strong incentive to assure a strong security baseline and continued familiarity, the lack of clear requirements for implementation or default security policies gives manufacturers broad leeway where other incentives, such as simple plug-and-play baselines with weak default security, may be favored. With approximately two-thirds of consumer devices maintaining default settings [46] or failing to keep up with system or security updates [181], internal devices' exposure to external threats becomes dependent on the router's design. Without a default security perimeter in place, once "secured" devices within a home

network would now rely on the consumer to either individually maintain each device or to implement a technical solution, such as detailed firewall rules, on their own.

In this chapter, we assess a variety of gateways from popular manufacturers in order to glean whether the lack of address scarcity under IPv6 is changing the incentive model that previously lead to common security baselines and default security policies being implemented. Our findings show large inconsistencies in the implementation of IPv6 default configurations, service exposures, and an overall lack of messaging to consumers about the baseline policy of a device. As a result, in cases where no default firewall is enabled, consumers may be unaware of the exposure to their devices while developers may have incorrectly assumed that a device's services are not exposed to the Internet. These results show a clear shift in baseline operation, demonstrating both the strong incentivization NAT provided for familiar security baselines and the need for new incentivization under IPv6 to ensure the same moving forward.

3.1 Introduction

For over twenty years, IPv4 network address translation (NAT) dictated a common operational template for customer edge (CE) routers across a diverse set of hardware manufacturers. Fueled by Internet growth and address scarcity rather than intended design, the ubiquitous usage of NAT, combined with RFC 1918 addressing, provides consumers and developers with a common behavioral standard [**?**, 193]. While unintentional, NAT meaningfully isolates devices inside the network from those outside it. This allows device manufacturers, and consumers by proxy, to benefit from automatic and default attack surface reduction.

In contrast, IPv6 provides enough address space that individual devices receive their own public, globally-routable addresses. This model eliminates the need for NAT and allows other devices on the Internet to communicate directly with devices in the home. The IETF provides little guidance or standard for firewall configurations [161, 28], allowing router manufacturers to implement filtering policies at their own discretion. With approximately two-thirds of consumer devices maintaining default settings [46] or failing to keep up with system or security updates [181], internal devices' exposure to external threats becomes dependent on the router's design. Without a default security perimeter in place, once "secured" devices within a home network now rely on the consumer to either individually maintain each device

or to implement a technical solution, such as detailed firewall rules, on their own.

In this chapter, we perform the first study of IPv6 CE routers to examine how manufacturers are implementing filtering and access control for IPv6 residential networks. We assess ten popular CE routers to evaluate their default firewall policies and the ability for consumers to implement custom rules. Our findings show inconsistency in the implementation of default configurations, overexposure of services, and an overall lack of messaging to consumers about the baseline policy of a device. As a result, in cases where no default firewall is enabled, consumers may be unaware of the exposure of their devices and developers may have incorrectly assumed that a device's services are not exposed to the Internet.

The remainder of this paper is structured as follows: In Section 3.2, we provide a short overview of IPv6 features, operation considerations and competing security paradigms. We then present our methodology for assessing IPv6 implementation in CE routers across a spectrum of features and configurations in Section 3.3 before presenting our results in Section 3.4. We discuss the necessity for a single device baseline standard and recommend consistent messaging in Section 3.5. Finally, we conclude in Section 3.6.

3.2 Background

Although functionally similar to IPv4, IPv6 provides a few small but impactful changes to the typical consumer network. In this section, we give a brief history of the transition from IPv4 to IPv6 before covering some key differences between the two protocols and their potential impact on consumers.

3.2.1 IPv4 NAT

NAT shaped the CE routing environment for two primary reasons: First, the scalability of NAT delayed the eventual address exhaustion of IPv4 in a period of explosive Internet growth and provided a simple path to connect significantly more devices to the Internet. Internet Service Providers (ISPs), who manage public address distribution in their networks, effectively required CE routers to support NAT by allocating exactly one public IP to each household gateway [66].

Second, the simplicity of NAT lowered the barrier for non-technical users to operate their own network. Home networks are often unmanaged or rely heavily on default configurations to meet the

needs of non-technical users [28, 193, 32]. By adopting NAT, CE routers were able to provide simple or automatic initialization that required minimal configuration beyond Service Set Identifier (SSID), Wi-Fi Protected Access (WPA) password, and any ISP-specific settings (such as a PPPoE username/password) [32]. Once established, a suite of protocols (UPnP, STUN, etc.) provide an interface for connected devices to negotiate with the router directly such that the user would rarely need to interact with the network [134, 192, 161]. NAT also removed the need to define and manage an ingress filtering policy, as the one public address is multiplexed for use by all internal hosts. The prevalence and ubiquity of NAT are now synonymous with the default-deny ingress policy that has become the de facto security model of CE networks, a policy that *is often the only ingress access control deployed.*

However, the motivation for the adoption of NAT in IPv4 is negated by a core feature in the design of IPv6: there is no longer an addressing shortage meaning we again have the ability to assign one or more addresses to each device. With this transition, inbound access controls are now discretionary; IPv6 allows CE networks to operate without the network perimeters and default access control necessitated by NAT.

While the IETF explicitly acknowledges that care should be taken in designing the baseline operation of CE routers, they avoid proposing default configurations due to a constructive tension between the desires for transparent end-to-end connectivity on the one hand, and the need to detect and prevent intrusion by unauthorized public Internet users on the other [193]. The strongest recommendation provided by the IETF is for manufactures to include a toggle to allow customers to choose between an open, unfiltered gateway where security is left to endpoint devices, or a closed perimeter approach, similar to NAT, where traffic is filtered and only allowed through careful exception [28, 193]. In the absence of efforts by manufacturers to provide standardization or documentation of the defaults that they implement, consumers are left to assess whether the security model that their network implements is sufficient.

3.2.2 IPv6 Reachability

A significant consideration in the adoption of IPv6 is the ability to uniquely address each device that joins the Internet. No longer defined by NAT architectures and private subnets, this addressing allows for every device to be globally *reachable*. Devices designed for the home environment often pose

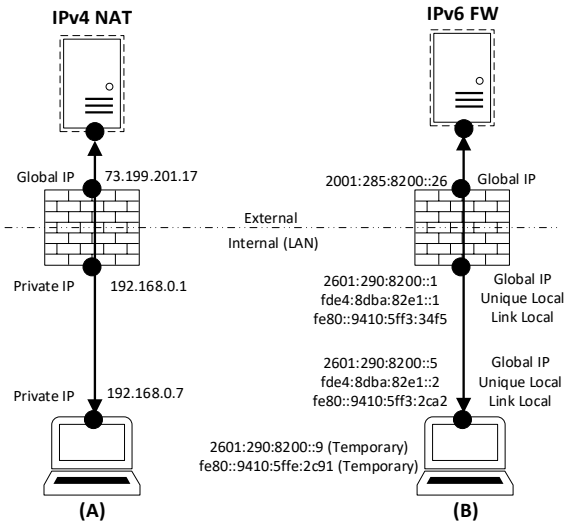

Figure 3.1: **IPv6 Service Validation Network Testing Layout** – IPv6 represents a fundamental shift in the addressing of local networks. In (A) NAT, computers follow a one-to-one mapping of local network private IP with a single globally-routable IP shared by many internal devices. In (B) IPv6, devices can have many addresses depending on operational scope. Additionally, IP addresses are unique and can be routed globally if it has the correct scope – a direct contrast to IPv4 NAT.

a serious risk when exposed to the open internet [94, 92, 11]. However, globally reachable does not automatically imply a device is globally *accessible*.

The IETF's RFCs give router manufacturers discretion for handling unsolicited inbound traffic in IPv6. The two basic options for default policies are:

- **Default Deny**: drop all unsolicited WAN-to-LAN inbound traffic. To permit inbound traffic, users can either manually add firewall exceptions or rely on protocols that allow exceptions to be negotiated directly with the router. This policy resembles the existing model of IPv4 networks instrumented with NAT and UPnP.
- **Default Permit**: allow unsolicited inbound WAN-to-LAN traffic. Devices are globally accessible, offloading the responsibility for filtering unwanted traffic to each individual device. The advantage of this model is that developers can easily design and deploy their Internet-capable devices without consideration for including and maintaining additional security mechanisms such as firewalls, hole punching mechanisms, or their associated user interface controls.

Whichever default policy is used, the mental model that a user employs must change from that of IPv4. If a user wishes to manually configure an exception to the ingress policy that their router

implements, the subtle difference between NAT and individually globally-addressed devices is significant. For example, individual devices in IPv6 can have more than one address assigned concurrently, and those addresses may be link-local or transient as demonstrated in Figure 3.1. In order to administer their IPv6 network in a manner equivalent to IPv4, users must understand technical details about IPv6 operation and firewall behavior. This is further complicated by the fact that the control interfaces provided by manufacturers and across devices have no common nomenclature or abstractions for configuration tasks. A study of enterprise IPv6 networks found that enterprise operators likewise have difficulty implementing appropriate controls in these networks [44]. These challenges should not imply that there is anything inherently wrong with IPv6 - the same model provided by IPv4 NAT can similarly be implemented in IPv6 [185] - but further demonstrate the need to provide a common expectation for baseline operation.

The flexibility of implementation among CE routing devices combined with globally reachable addressing creates a potential issue: unlike IPv4 networks where the de facto model is effectively required, in IPv6 CE routers are free to expose all internal endpoints. Furthermore, as devices transition from IPv4 to IPv6, this exposure could occur without any communication to the end user as they attempt to administrate their network. Because inbound access control implementation is left to the discretion of manufacturers, we suspect that there is variance among implementations. In the next section, we describe our methodology for evaluating a set of off-the-shelf CE routers to assess how IPv6 access control is implemented in practice.

3.3 Methodology

Our study aims to measure the security implementation of consumer grade gateways and the configuration options that they provide for IPv6. In this section, we describe our methodology for selecting and evaluating these routers.

3.3.1 Router Selection and Network Configuration

In order to choose routers that are representative of those deployed in real networks, we rely on the work of Kumar *et al.*, who provide insight into the most commonly used global gateways by manufacturer and region [94]. Out of 4.8K router vendors globally, we selected 12 routers which covered 25.2% of the

most commonly deployed global brands. Only routers that specifically mention compatibility with IPv6 were chosen for our comparison. We were unable to find any routers that advertise or provide messaging about filtering policies.

To evaluate the potential differences within a manufacturer we include multiple Linksys (`EA3500`, and `EA6350`) routers. Two of the selected routers (the `Tenda AC18` and the `Wavlink Aerial G2`) were excluded because they did not actually support IPv6 upon arrival. The remaining ten devices used in our assessment are shown in Section 3.4, Table 3.1.

Our architecture consists of four key elements marked with letters in Figure 3.2. Two vantages were established to assess traffic flows: an external host located on a public cloud provider (A) scanning across a public ISP toward the firewall (B) or internal host (C), and an internal vantage (D) which conducted the same scans focused outbound (with the exception of targeting an external host due to the ubiquitous outbound permit policy of the firewalls).

All devices sending and receiving probes associated with scans were under our control at all times and at no time did we perform any scanning or analysis of public or private systems outside of our controlled scope. This architecture allowed us to pass traffic across the public internet via local consumer grade ISPs and through the assessed routers from different vantages to analyze real-world operational modes.

3.3.2 Evaluation Methodology

In order to allow unsolicited inbound connections (e.g., peer-to-peer connections), IPv4 routers must provide the ability to *port forward*; the router establishes a list of port numbers and destination (internal) addresses. When a packet is received on the public interface at a port in the list, the router bypasses any NAT lookup and immediately rewrites the destination address and forwards the packet internally. Forwarding is common in IPv4; devices rely on the UPnP and NAT-PMP protocols to automate the setup of forwarding rules. Without these protocols, users would need to manually create such rules, a technical task requiring knowledge of IP addresses and TCP/UDP ports.

Forwarding is effectively meaningless in IPv6 without NAT as devices can be addressed directly. Instead, routers must provide a mechanism to create firewall exceptions if a firewall is implemented.

While these rules can be as simple as port forwarding rules (e.g., a destination IP and a port number), how they are implemented and the options available to users may vary. We evaluate the following basic characteristics of each router:

- **Default IPv6.** We first check if each router supports IPv6 and whether it enables that support by default. When IPv6 is enabled by default, IPv6-capable devices on the internal network automatically request addresses. Default IPv6 support requires that the upstream ISP also supports IPv6. It is notable that router support for IPv6 and default enable state can be changed in a firmware update pushed remotely by the manufacturer, and ISPs can (and do) add support for IPv6 without notifying consumers. Therefore, *devices in the home environment can transition to IPv6 overnight without the user's knowledge.*
- **Firewall Present.** Next, we evaluate whether or not the device implements a firewall. In cases where a firewall is not present, the device will pass all traffic to internal hosts.
- **Firewall Enabled.** If a firewall is present, we evaluate whether or not it is enabled (i.e., filtering) by default.
- **One-Click Open.** While RFC 7084 refrains from proposing a default IPv6 ingress filter policy for consumer gateways, it advises that gateways implement a single button to toggle all firewall ingress filtering [161]. We evaluate whether or not the device includes this functionality.
- **Security Warning.** When the One-Click Open option is used, we evaluate if there is any warning or communication to the user about the danger of disabling the firewall.
- **Rule Generation.** We evaluate whether each device includes the ability to create exceptions to the default firewall policy. Such rules may be necessary for allowing specific services or applications to function in the presence of a firewall. Because we are comparing to existing functionality in IPv4 networks, we specifically exclude examining more expressive firewall capabilities than IP/device/port tuples.
- **IP Specification.** We evaluate whether or not rule creation specifies an individual IP as the destination.
- **Device Specification.** As IPv6 devices are often assigned multiple addresses (in some cases,

one per application), creating a rule may be complicated by device/address identification. We evaluate whether rules can be created by specifying a device (e.g., by MAC address or another identifier) rather than a specific IP address.

- **IPv6 UPnP Support.** Finally, we evaluate the router's capability to offer *automatic* rule generation. Devices on the local network can use UPnP to create firewall rules programmatically if the router offers this capability.

Since routers do not explicitly advertise their firewall policies, we conduct a series of black-box scans in order to establish the default filtering model, firewall filtering policies, and hosted router services. We designed and built a custom traffic monitor on the internal host to ensure accurate collection of packets arriving through the firewall. During a scan, this monitor would listen for and record inbound IPv6 traffic with a timestamp, arrival port, protocol and scanning source IP. We reconciled the packets received with packets sent from the scanner to filter unwanted traffic and verify correct operation.

Scans were conducted using Nmap against the most common 1,000 TCP and UDP ports (as defined by the scanner). This scope was chosen due to interest in exposure of the most common ports and scan duration considerations. A complete assessment of each CE router involved nine total scans from two sources, each conducted with the firewall on and off as shown in Figure 3.2: First, scan (1) is conducted from the external vantage to the internal host establishing the inbound filtering strategy of the firewall. Scan (2) probes the external router interface from the external vantage to identify open ports and exposed services; (3) repeats this scan on the internal interface to determine if this traditionally concealed interface is exposed under IPv6. For each interface, we conduct a banner scan against exposed ports (4 and 5). This process is repeated from the internal vantage first targeting the exposed services on each router interface (6 and 7) before conducting the same banner grab on exposed services (8 and 9). The combination of sources and targets allowed complete measurement of IPv6 filtering policies, exposure, and default operational model of the CE router. These results were then compared with our evaluation of basic router characteristics to complete a holistic router assessment, presented in Section 4: Results.

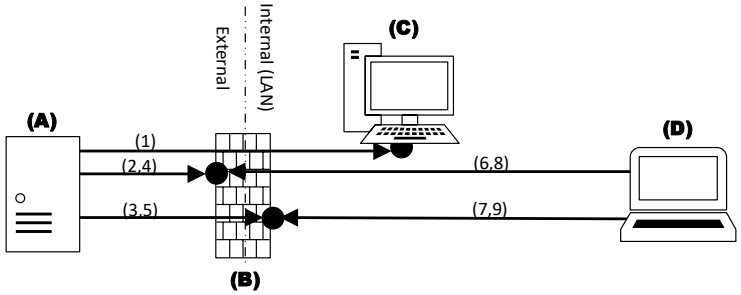

Figure 3.2: **IPv6 Router Scanning Protocol** – To fully evaluate the security policy of each router we scan from two vantage points (A) and (D) against three targets: (C) an internal host, and (B) the firewall internal and external interfaces. In total, we conducted 9 unique scans for each router.

Table 3.1: **Selected Router Baseline Configuration Metrics** – This table displays the heterogeneous nature of management options and default configurations among the devices evaluated. Bolded device names indicate that the router implements a default-permit firewall policy and IPv6 is enabled by default. Configuration options for unsupported features are marked with dashes. No device used IPv6 NAT.

Device	Brand	Firmware Version	Default IPv6	Firewall Present	Firewall Enabled	One-Click Open	Security Warning	Rule Generation	IP Specification	Device Specification	IPv6 UPnP Support
Amazon Eero	Amazon	Eero OS 3.15.2-1	○	●	●	○	○	●	○	●	○
AmpliFi Gamer's Edition	Ubiquiti	v3.3.0	○	●	●	●	●	○	–	–	●
Cisco DPC3941T XB3	Cisco	2.3.10.13_5.5.0.5	●	●	●	●	○	○	–	–	○
Google Nest (2nd Gen)	Google	12371.71.11	○	●	●	○	○	●	○	●	○
Linksys EA3500	Linksys	1.1.40.162464	●	●	●	●	○	●	●	○	○
Linksys EA6350 AC1200	Linksys	3.1.10.191322	●	●	●	●	○	●	●	○	○
Motorola MR2600	Motorola	1.0.10	●	●	○	○	○	●	●	○	○
Nighthawk X4 R7000	Linksys	1.0.0.124	○	●	●	●	○	○	–	–	○
Surfboard SBG10 DOCSIS 3.0	Arris	9.1.103AA72	●	●	●	●	○	●	●	○	○
TP-Link AC1750 v2	TP-Link	180114	●	○	–	–	○	○	–	–	○

3.4 Results

In this section, we present the results of our experiments for each of the CE routers. In general, we find CE routers with IPv6 capability have little commonality of security implementation across manufacturers.

3.4.1 Operational Defaults

Table 3.1 presents an overview of our findings showing a wide variance in default operation, security, and user control. Eight of the ten routers assessed have an enabled default firewall policy (i.e., default-deny) for IPv6 while the remaining two devices (`TP-Link AC1750` and `Motorola MR2600`[1]) do not have a default firewall (i.e., default-permit). Neither of these two devices communicates this design decision to the consumer. At the time of writing, the `TP-Link AC1750` is Amazon US's top-selling router [9] and TP-link is the top global provider, accounting for 15.9% of all deployed devices [94], suggesting that the default permit model may be commonly deployed.

Five of these eight default deny devices further provide a "One-Click Open" option for opening the network to inbound connections. This option immediately transitions the network to a default permit model allowing all ingress traffic through to the internal hosts. The effect that this has on ingress filtering can be seen in Figure 3.3 in the Appendix. Only one of the ten devices evaluated provides an explicit warning to the user before allowing the firewall to be disabled using this feature. Users with minimal technical knowledge who are accustomed to a default closed model from IPv4 NAT may be unaware of the additional exposure this option creates.

Two routers, the `Motorola MR2600` and `TP-Link AC1750`, enable IPv6 routing by default with a default permit firewall. This combination of configuration settings exposes all IPv6-capable devices to the wider Internet by default. While the `Motorola MR2600` allows consumers to *optionally* enable the firewall, the user must be aware of the current state and possess the technical capability to do so. Worse, the `TP-Link` router only provides the ability to disable IPv6 and has no capability to enable any filtering.

3.4.2 Firewall Policies and Pinholing

We find a spectrum of firewall management options offered to the consumers ranging from subscription model services for packet inspection and filtering, to singular on/off toggles, to complete lack of firewall configuration for IPv6. Depending on the router, modifying the configuration can be accomplished

[1] *Responsible Disclosure* Given the severity of enabling IPv6 support by default and a default-permit posture, we disclosed our findings to both Motorola and TP-Link in August 2020. In November 2020, Motorola issued a public patch to correct the issue. TP-Link did not respond to our disclosure.

through a smartphone application or a locally hosted web portal, with a few devices supporting both.

For routers that provide an interface to create exceptions to the default firewall filtering policy (pinholes), we found that two out of six connect those rules to the device MAC address. We verified that in these cases, traffic destined for *any* associated address for the device is forwarded. The other four out of six routers allow users to provide a single, static address that the rule applies to; the rules are not updated if the device migrates or is assigned additional IPv6 addresses over time.

Of the routers that do not support IPv6 pinholing, only the `TP-Link AC1750` provides no ability to configure the firewall aside from disabling IPv6 (because it does not have such a firewall). For the remaining three routers, `Cisco DPC3941T XB3` also provides several options of choosing what kind of traffic is blocked besides the "One-Click Open" option, while for `Ubiquiti AmpliFi` and `Netgear Nighthawk`, *One-Click Open is the only method available for users to control the firewall.* As an example, the `Ubiquiti AmpliFi` provides users with minimal control over IPv4 policies through port-forwarding controls, but the management interface lacks an equivalent ability to create pinholes in IPv6. Ubiquiti notes this on their official FAQ: "AmpliFi does not support editing firewall configurations, and cannot be disabled unless you place the router in bridge mode" [127]. Contrary to this statement, they do allow automated modification of firewall rules through the embedded UPnP `WANIPv6FirewallControl:1` device template. For manual control, the web interface instead offers an "Allow all incoming IPv6 connections" as the only actionable solution for non-technical users.

3.4.3 Router Scanning

We find that when CE routers are globally accessible a majority of them expose open services to the Internet as shown by Table 3.2. Whether the firewalls are disabled manually or by default, six routers do not employ rules to restrict access to local network services from the global Internet. We found that services (e.g., SMTP, HTTP, and SMB) available on internal router interfaces were also offered on the external interfaces as well as the link local address on these devices. Interestingly, this indicates that the manufacturers are configuring their internal services to listen on all interfaces; when the firewall is off, these services are no longer protected. It is unclear if this is an oversight or expected operation.

Table 3.2: **Externally Exposed Services for Assessed Routers** – This table lists the IPv6 services and open TCP ports that are exposed by each device with the firewall either enabled or disabled for the routers that support such an option. Ports in bold indicate that a service responded with a banner. We document the services associated with the address from the router's external interface. Most routers have a separate address assigned to their internal interface from their allocated subnet, though we find that the exposed services are typically the same between the two.

Device	Default FW	FW Enabled	FW Disabled
Amazon Eero	●	–	No Disable Option
AmpliFi Gamer's Edition	●	–	–
Cisco DPC3941T XB3	●	–	–
Google Nest (2nd Gen)	●	–	No Disable Option
Linksys EA3500	●	–	25, 53, **80**, 135, 139, 443, 445, **2601**, 1080, **10000**
Linksys EA6350 AC1200	●	–	25, 53, **80**, 135, 139, **443**, 445, **2601**, 1080, **10000**
Motorola MR2600	○	25, 135, 139, 445, 1080	25, 135, 139, 445, 1080
Nighthawk X4 R7000	●	–	25, 43, 80, 135, 139, 443, 445, 548, 1080, **2601**
Surfboard SBG10 DOCSIS 3.0	●	–	25, 80, 135, 139, 443, 445, 1080
TP-Link AC1750 v2	○	No Enable Option	**22**, 25, 135, 139, 445,1080

We discovered two exceptional implementations: First, the `Motorola MR2600` maintains a small subset of exposed open ports on its external interface even with the firewall enabled. Second, the `TP-Link AC1750` maintains an outdated version of Dropbear SSH despite the public availability of a CVE describing a remote code execution vulnerability [113].It is notable that, of the routers that expose ports in any firewall configuration, there appear to be a common set of ports that are open, but provide no banner. We hypothesize that these ports are associated with common services that each router provides but does not enable by default, though the ports remain open. For example, multiple routers advertise the ability to set up local storage sharing, likely using SMB on port 445. Though we did not exercise this functionality, the exposure of these ports suggests that if a client were to enable these features they would also be accessible to the wider Internet over IPv6. The default states and mix of services available provide enough unique scan data to individually identify the device manufacturer; six of the ten routers we obtain have uniquely identifying features. As a result, we believe it may be possible to fingerprint routers through probing open IPv6 ports and services, though we leave this to future work.

To summarize, our work shows that there is little standardization among the routers evaluated in this work around the security or operational functionality provided for IPv6 CE networks. This is in direct contrast to IPv4 where devices and services are not exposed. While NAT was not designed as a security framework, the deny-all, permit by exception ingress policy serves as an invariant for consumer routing devices and is noted as such within RFCs when debating the default recommendations of CE routers [193, 28, 161]. We see this argument manifest in the inconsistency between device implementations; the default policies maintained by devices put real users and systems at risk.

3.5 Discussion

The current lack of a clear operational model for IPv6 within consumer gateways fails to learn one of the key lessons taught by IPv4 and NAT – that the Internet will leave standardization behind if there is demand incentivizing the delivery of a capability.

As discussed previously, the IETF has refrained from requiring either an open, end-to-end approach, or a more familiar closed model with a well defined perimeter similar to NAT. This lack of formal requirement has lead manufacturers to implement IPv6 disparately. The IETF cites this lack of formal definition as "constructive differences" within the community on desired approaches [28]. We argue this is a failure on the part of the IETF to learn from the lessons of IPv4 and NAT which puts (more often than not) non-technical consumers at the mercy of a non-heterogeneous IPv6 deployment.

3.5.1 Need for a Single IPv6 Operational Baseline

What is clear from our review, at present, is gateways operate IPv6 with no clear security baseline. In many cases we find the default policies, and mechanisms by which to adjust, provide significant exposure to the consumer. With many end devices prioritizing IPv6 use, it is likely that a consumer may already be operating IPv6 without their knowledge. Alternatively, an ISP may chose to enable IPv6 routing resulting in a customer having a stateful filter with NAT one day and potentially nothing the next. While the IETF has provided working recommendations in support of a default standard, many of the identified requirements are optional or remain open to interpretation [193]. This lack of precise

definition echoes the approach used to define NAT, and with it, the challenges that ambiguity enables.

While NAT provided an assured security baseline through a default-deny filtering policy, the same assurances are not present under IPv6. In 2019, the IETF noted this challenge in their consideration of security recommendations, stating, "In new IPv6 deployments it has been common to see IPv6 traffic enabled but none of the typical access control mechanisms enabled for IPv6 device access [138]". Others have found that IPv6 devices are twice as accessible compared to IPv4 and further exhibit unique vendor response behaviors, similar to the differences presented from a lack of definition within NAT [44]. The end result is more in line with "assured exposure" instead of a secure default for the non-involved user. To address, clear standards that precisely define baseline configuration and security policy could serve as strong incentivization for compliance, where the risk of non-compliance could result in potentially costly lawsuits [**?**].

3.5.2 IPv6 is not IPv4

While it is easy to compare and assume similar operational characteristics between the two protocols, this is dangerous in practice. Within the home gateways we assessed, there is clear demonstration of manufacturers re-implementing IPv4 policy mechanisms for IPv6 while failing to account for key operational differences between the two protocols. Systems that filter IPv4 hosts based on IP do so by correlating single addresses to a single host. Transferring this same reasoning to IPv6 does not account for shifts towards multi-homing and separate operational scopes of addressing or ephemeral use, often resulting in a security mechanism being present for use but not actually providing the intended control.

In a similar consideration, absent a stateful filtering policy similar to NAT, many services intended for the local or "trusted" side of the network were broadly exposed to the open Internet when operating under the open model of security (either by default or by actively disabling the firewall). While NAT provided a default stateful filtering policy, it appears that manufacturers are incorrectly relying on this presence in IPv6 as many of the devices we assessed exposed local services when the firewall was disabled. This problem is not relegated to consumer grade network equipment as studies have shown the same challenges present across enterprise deployments [99, 44]. Implementation of a IPv6 stateful filter similar to NAT would help address many of these challenges.

3.5.3 Consumer Involvement

While consumers could likely forgo implementing their own security policies with NAT, this hands-off approach carries significant risk with IPv6. It remains unclear whether or not a non-technical consumer should have any expectation to participate at all, absent an individual desire to provide a more refined policy. In many of our assessed gateways, the opposite is true, demanding of the user both an active involvement to secure and a technical understanding to do so precisely. This is a significant paradigm shift that is neither communicated to the user (via packaging, setup, or broader industry communications) nor are they given the tools to do so effectively.

While the IETF does acknowledge that the expectation and role of the consumer to likely be limited, there is little alignment to these principles being demonstrated by the gateways themselves. In reviewing both the historical and present challenges surrounding these devices, it is clear that consumer security has *never* been a leading design consideration. In cases where manufacturers have presented mechanisms to aid or abstract user involvement, further exposure is commonplace. With UPnP alone, billions of routers have been exposed through underlying security flaws [20, 58]. Despite ten years of efforts advising consumers to disable this feature, UPnP remains an on-going challenge [160]. Efforts to address this lack of consumer involvement through automatic security updates still presents an incomplete solution, as many updates can lag exposures by months or years [68].

Alternatively, guiding users towards a secure configuration, such as through a guided setup process, could likely improve outcomes. Here users are incentivized to complete (they want their network up and running) which manufacturers could leverage to establish stronger security actions through a setup and configuration process. This forces consumers to make decisions about their security rather than leaving it to manufacturers to guess. While approaches to guide users towards secure configurations are not new [95], precisely balancing configuration flexibility, terminology, and selection of initial tasks can be challenging over a broad user base of skill and need.

3.5.4 IoT Security Considerations

Of greatest concern is the effect that these unclear default policies and control mechanisms will have on the devices within a consumer network. In particular, IoT and Smart Home devices present a unique challenge. Low cost design and hardware limitations prevent many of these devices from providing feature rich security mechanisms [104, 70, 162]. As a result, many devices overwhelmingly rely on simple local authentication methods as the only means for access control[70]. However, this does not prevent reachability of a device or limit its behaviors on a network.

Further challenging security is a preference to utilize IPv6 or 6LoWPAN for network connectivity [183]. While these protocols do provide for organic security measures, such as encryption and authentication, many devices cannot support the computational overhead introduced [162]. While research into providing lightweight security protocols is an active and open problem for the IoT community, the immediate and easy solution is to simply forgo these measures and rely on perimeter defense mechanisms [70]. The result is both an increased need and expectation for consumer gateways to precisely provide default security assurances and mechanisms by which to adjust.

3.5.5 Open-ended Design

The CE environment provides a unique challenge in balancing device capability against user ability and need. This work demonstrates that the shift to IPv6 removes the consistency of one of its most crucial layers of defense: homogeneity in router operation. Without a safe default policy, consumers must rely on the security of each of their endpoint devices, which can be difficult to ensure, especially in CE environments where device maintenance is not guaranteed. We recognize that many of these problems are not caused by or unique to IPv6 consumer networks, but we note that unclear IPv6 implementation strategies exacerbate these issues by offloading responsibility for securing and configuring the network to consumers.

We see in our assessment a struggle to shape and define what exactly is the right amount of control without under-offering or overwhelming targeted consumer demographics. This has left router manufacturers to determine what are the correct abstractions and implementations, and how to commu-

nicate these clearly to a wide demographic of users. Accordingly, we believe that addressing the general inconsistency is the most direct path to securing CE networks in IPv6.

3.5.6 Recommendations

There are multiple parties involved in CE environments each of which have different motivations, incentives, and risk factors, but it is important that the design of CE networks prioritizes the wholesale security of consumer data and devices. We structure our recommendations around the following principles:

- The default operation mode should be secure, and the bulk of network configuration should be moved from consumers to developers.
- Configuration options should be consistent and only as permissive as necessary.
- Configuration pitfalls should have confirmation warnings that ensure users understand the risks associated with the changes they are making (**e.g.,** making devices globally accessible).
- Documentation should share abstractions and language across manufacturers and be as minimally complex as feasible.

It is important to present a clear, consistent threat model to consumers whose ability and understanding often lags that of developers, to avoid oversight on responsibility for securing devices connected to home networks. This is the responsibility of both standardization bodies and the CE router industry as a whole. We strongly recommend the following defaults:

3.5.6.1 Standardization

We recommend that CE routers universally standardize around a default ingress filtering policy that denies incoming traffic. We further recommend manufacturers remove or restrict the "one-click open" option on CE routers as home users are likely to unknowingly expose their whole network, violating the security principle of least privilege. If this is a required functionality, routers should warn users (and/or suggest to use IPv6 pinholing) before allowing them to use this option.

For manual exceptions we recommend that manufacturers implement both device and IP based rules and develop a consistent vocabulary for describing them. Providing users with the resources

to understand when each option is preferable will require that the language used to describe IPv6 configuration options is consistent across manufacturers.

3.5.6.2 Documentation

It is irrelevant what standards require if manufacturers ignore them or if parties involved fail to understand their importance or the importance of their abstractions. Fostering consumer and developer understanding of IPv6 security can create pressure on manufacturers to adhere to standards and promote transparency ahead of purchase. Establishing consistent language and abstractions for describing the security mechanisms of IPv6 networks is the first step.

Currently manufacturers of customer edge routers highlight IPv6 as an enhanced feature in their product marketing, though we found no instance of educating users about IPv6 or describing its security implications. Instead, phrases such as "provides infinite addresses for more devices", "best possible experience", and "simplifies the router's tasks" are offered as slogans to encourage user commitment [111, 100]. These approaches are problematic. This hides a transparent shift in the security model of home networks that consumers cannot be expected to inherently understand on their own.

Morgner *et al.* present one possible solution of offering device label standards similar to nutrition labels on food [119]. Here, the authors focused on manufacturer guarantees for duration of product support and timeliness of updates in a standardized label. We argue to take this concept further with a holistic approach to additional aspects of security such as default configuration, control mechanisms, and 3rd party certifications. Requirements for labelling standards incentivize manufacturers to provide and document security features necessary for consumers to have a functional understanding of their network posture at purchase.

3.5.7 Future Work

While this work discusses at length the "One-Click Open" option, we have not conducted a formal user experience study to confirm that users will rely on this option to achieve simple routing changes in their IPv6 networks as a first choice. A proper study of the UX/UI design involved in home network

security would be informative and could provide developers with a better understanding of consumer needs and approaches to IPv6 security.

While we use this work to gauge the scope of current security policies of IPv6 CE routers, a large scale examination of router IPv6 firewall behavior is required to better understand the breadth of the impact that the transition from IPv4 to IPv6 has on CE routing. Specifically, a tool assisting clients to better understand the defaults that their network implements could prove a strong contribution towards this result. Similar large scale studies of IoT and smart devices operating in IPv6 environments are reserved for future efforts as well.

3.6 Conclusion

In IPv4 networks, the use of NAT afforded a ubiquitous, de facto default-deny security posture. The growing deployment of IPv6, which eliminates address scarcity, no longer requires NAT. In the absence of strong guidance for how router manufacturers should implement filtering, we examined a diverse set of routers to measure real-world implementations. We find that the access control models and controls implemented to manage these networks are coarse and contain unsafe defaults that likely expose devices on the network – often without warning to the consumer. The result is a systemic, demonstrable failure among all parties to agree upon, implement and communicate consistent security policies.

Given that many of the issues under IPv6 have similar ties to the early challenges of NAT, it is clear that we have not learned the right lessons and that incentives to provide stronger security are missing. While previous academic work has recommended a number of approaches, such as consumer security labels, to encourage manufacturer compliance [157, 84, 51], no manufacturer or standards agency has taken it upon themselves to introduce or follow such a recommendation. While this approach could be beneficial for both user awareness and involvement in security, the present opt-in approach of such an idea fails to provide strong incentivization for participation.

In light of the fact that the current self-governing model for security is not working, stronger incentives are necessary to enact the broad changes necessary to prioritize security within these devices. Here, defining clear standards could go a long way towards enabling clear security baselines as manufacturers

would have a defined requirement to follow. Not doing so could potentially expose them to lawsuits, serving as strong incentives for compliance.

If stronger incentives are needed, leveraging regulations has shown a clear and measurable effect on establishing stronger security outcomes, as demonstrated by policies in the payment card and banking industries [170, 184]. One unique regulation for device security was recently enacted by California which required internet connected devices to have both a unique default password and requirement to change this default password on initial setup [15]. Uniquely, there appears to be no cost for non-compliance, making it unclear what effect this law will ultimately have towards incentivizing security outcomes.

It is important to note, compliance is a business decision to balance costs and risk. If there is no cost associated with a risk, there is no incentive for change. If we are to see better security outcomes for home gateways, we need to have larger costs associated with non-compliance. The strongest incentives to force change are standardization and regulation, assuming that costs for non-compliance are present and enforced. Within the home gateway market in particular is a lack of regulation, oversight, or direction to both define and hold accountable manufacturers and their requirements for security. The self-governing model used at present by the internet community has not been enough to ensure the broad application of security, necessitating a change in approach.

Alternatively, we can build incentives into a design, thereby helping to encourage adoption. In the next chapter, we demonstrate one such approach using a global database of routing information, which we use to demonstrate how system administrators could find initial value while helping solve challenges that they have at present. While the next chapter changes focus towards broader Internet security challenges, the core concept of building incentives into solutions as a means to enable security is broadly applicable across all environments.

3.7 Appendix

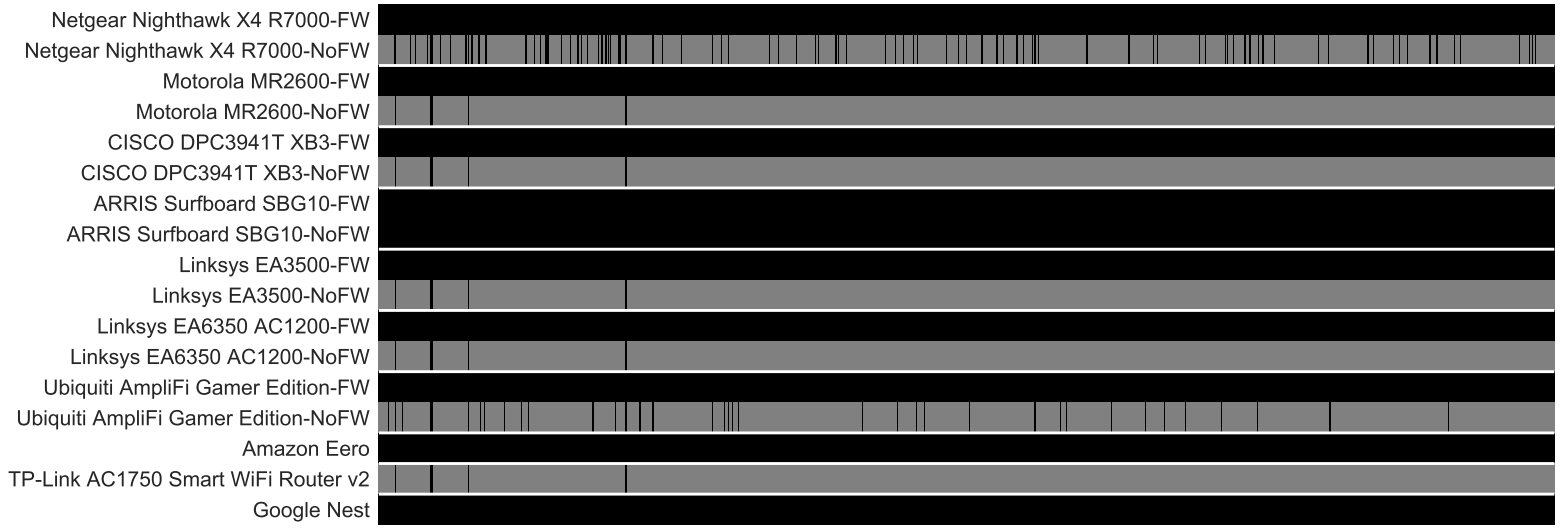

Figure 3.3: **Firewall Ingress Policies (TCP)** – We use Nmap to scan the most common 1000 TCP ports on an internal host from an external vantage point. For each packet the host receives we mark the associated port GREY. Conversely, if the firewall drops the packet or the packet fails to reach the host due to network failure the associated port is marked BLACK. For routers that have an optional firewall we include a scan in both states indicated by FW or NoFW.

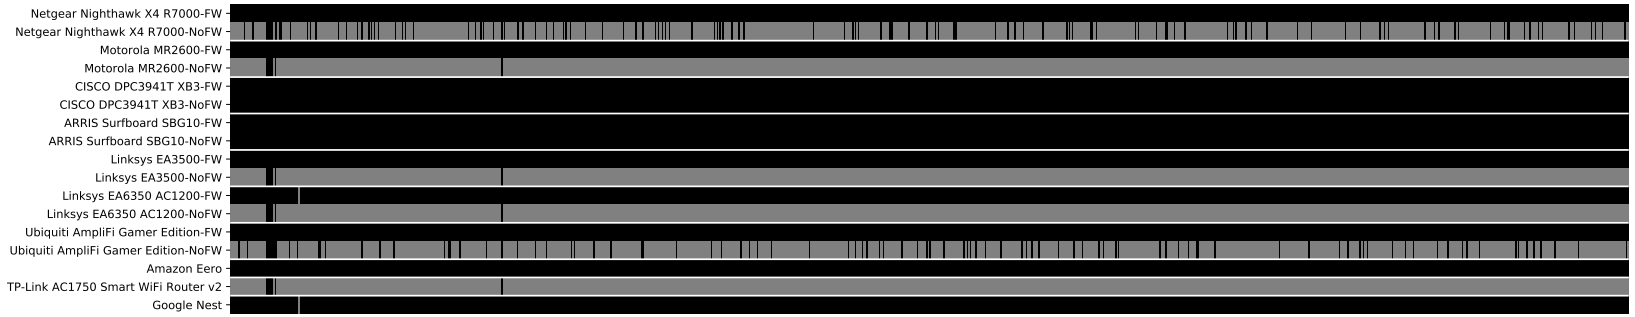

Figure 3.4: **Firewall Ingress Policies (UDP)** – We use Nmap to scan the most common 1000 UDP ports on an internal host from an external vantage point. For each packet the host receives we mark the associated port GREY. Conversely, if the firewall drops the packet or the packet fails to reach the host due to network failure the associated port is marked BLACK. For routers that have an optional firewall we include a scan in both states indicated by FW or NoFW.

Chapter 4

Building Incentives into Security for Better Adoption

As demonstrated with our prior findings, the present incentives to ensure a strong security environment are not enough. Little has been done to move the bar forward with regards to home gateway security over the last two decades of effort. Shifting focus to the protocols that underlay our networks, we find similar challenges. How to best secure BGP is a topic nearly as old as the protocol itself. Despite the desire to fix underlying security shortcomings, present solutions, such as RPKI, have struggled to achieve broad adoption. Underpinning this slow acceptance is the unclear value RPKI directly offers to a network provider in a low or partial deployment environment, where initial costs to implement and maintain do little to incentivize broad adoption.

Instead of designing clean slate approaches to BGP security, which often introduce complexity and cost, we should develop solutions that both incentivize and enable a provider to better perform functions around their core business. We can then leverage the adoption of these solutions to better introduce security mechanisms, overcoming barriers for deployment.

To demonstrate this concept, we propose a global routing database that network providers could primarily leverage to support management and troubleshooting of their own networks. Utilizing this database, we demonstrate how broadly democratizing network data can be beneficial to a provider and their business objectives. We then show how security approaches, similar to RPKI and BGPsec, could easily be adopted to this infrastructure. From this design, opportunities for new network paradigms can easily be created, allowing providers to leverage network data more broadly in the use of business objectives and routing security.

4.1 Introduction

Designed over thirty years ago, the Border Gateway Protocol (BGP) remains an ubiquitous and necessary protocol to support routing across the Internet. However, BGP is not without issue. Security challenges have been well documented in literature which include potential for human error [105, 27], data falsification[59, 105], protocol manipulation [115], and data misuse [168]. The end result is a global reliance on a system that assumes trust in operators to "do the right thing", despite regular occurrences to the contrary, whether intentional or not [159].

To address these shortcomings, a number of approaches to strengthen BGP security have been proposed. Methods to validate route origination [96, 187], provide assured path traversal [87, 97], and general best practices to limit propagation of bad routes [86, 24] readily exist. The use of blockchain is further proposed as a mechanism to support security based on distributed trust [17, 195].

One thing that stands out about each proposal is that they largely ignore incentives for deployment. A network provider deploying Resource Public Key Infrastructure (RPKI), for example, would introduce a number of new complexities to their environment imposing both technical and business costs. At the same time, the return on that investment is ambiguous at best, requiring broad participation by the internet community before benefits of RPKI are fully realized. Worse, attackers already have alternatives to circumvent [199]. The end result, unsurprisingly, is stagnant adoption. Ten years after standardization by the IETF fewer than 50% of autonomous systems have implemented RPKI, covering only 6.5% of Internet users [52].

The network research community needs to flip this problem around—we need to understand that network providers are a business first and build security solutions around that fact.

From a business perspective, solutions that ease management or troubleshooting, enhance business value, or enable performance efficiencies all serve as strong incentives for adoption. If we can align the components of a security solution to these incentives then the likelihood of adoption will be much greater, and maybe the enduring dream of BGP security can be realized.

To demonstrate this approach, we present a real-time Internet routing data-base to serve as a

foundational building block, which offers increased visibility for network administrators to manage, troubleshoot, and leverage network insights for business actions – key incentives for a business. In fact, providers likely already leverage similar solutions within their local networks. However, by centralizing and storing network data more broadly, there is potential to provide further value while also enabling greater security innovation. A global database already underlies RPKI, so a similar security approach would be on par while providing for a more flexible and evolvable design.

This sounds like both a simple and an impossible solution simultaneously. 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's proposing a central solution that sounds like it needs broad community participation. However, our proposal is centered on the concept that broad adoption should not be required for a solution to be beneficial. Instead, we should view security through the lens of whether it brings value to each provider, even if they are the only ones to adopt – we believe a global routing database does.

Utilizing the CAIDA AS-relationship dataset, we simulated a network of 53 ASNs leveraging Service Level Agreement and network transit cost data, stored within our global database, to enable *business optimal* topology outcomes. By including business requirements into the network decision making process, we show provider savings of up to 10% on overall incurred transit costs by leveraging broader types of data for route selection, demonstrating an initial incentive for providers to adopt a global database solution. Using the same approach, we further show a sliding scale of security, aligned to unique provider or customer requirements (which could introduce new product or pricing tiers), could be leveraged to select 100% compliant routes without requiring broad adoption rates, incentivizing early adopters. We leverage the results of these simulations to conclude with initial parameters necessary to achieve a global network database deployment.

The remainder of this chapter is organized as follows: In Section 4.2 we first consider the business of a network provider, the likely incentives they would consider in choosing to adopt a solution, and show how present approaches to BGP security provide little alignment to these incentives. We then introduce the architecture of our proposed solution (Sec. 4.4) centered around a global database of network data

that a providers can leverage for a variety of management, business, and security needs. We further demonstrate how security solutions, or other new paradigms valuable to a provider, could be extended from this centralized architecture through a series of experiments which we use to define a final global database design in (Sec. 4.5). We conclude with observations on challenges, opportunities, and directions towards realizing an incentivized approach to security.

4.2 Motivation

Motivating our work is a realization that present approaches to securing BGP, while addressing critical security flaws, rarely see implementation or acceptance in practice. In assessing both RPKI and BGPsec, the current approaches to providing route origin validation and path validation, we find designs that impose up-front costs to the administrator to implement while simultaneously delaying the intended value until broad community participation is realized. For example, BGPsec requires 100% participation along a route in order to establish an unbroken chain of signatures, a tall order for an environment defined by distributed participation and de-centralized control [97].

RPKI presents a similar challenge, relying on broad participation to establish published route origin authorization records (ROAs) by which adopters of RPKI must validate against [19]. If few providers employ, or deployment stagnates, localized deployments will do little to protect against secondary exposures resulting from an upstream provider falling victim to a hijacked route [7].

Beyond altruistic goals for internet security, there is little incentive for a provider to adopt RPKI at present. As an example, most providers establish value based on quantity of routes they can present to customers. By rejecting invalid routes,which could occur for a multitude of reasons, not always malicious, RPKI would seem to run counter to the core requirements of a business. Further, mis-configured ROAs can have an immediate effect on network availability, imposing unneccessary risk to an environment bound by service level agreements(SLAs) and performance. Incidents like this are not uncommon [74].

In either case, the presented solutions are static in their design, intended to address a singular functionality surrounding BGP security. As networks continue to evolve, it is unclear whether or not these approaches will even be required in the future or if a deployment at present will bring the intended

benefits if adoption continues to stagnate - strong disincentives for a business.

Assuming we are able to fix the shortcomings of RPKI overnight, adoption by administrators will still take years, if not longer, without alignment of value and incentivization for those who will be required to implement (network operators). Given these challenges, our work focuses on designing an approach where incentivization is central to security design while further providing a platform for evolution to easily occur- features not present in current BGP security approaches. Additionally, value should be realized, even if a single entity is the only one to implement. By prioritizing these design goals in-line with the community that will be the ones required to implement, we can help to remove the challenges creating stagnation in adoption that present approaches impose.

4.2.1 Aligning Solutions to fit Business Needs First

First and foremost, network providers are commercial entities whose *primary* product is to provide connectivity services to individuals, businesses, and other organizations at the lowest cost, enabling the business to turn a profit or maintain a competitive edge. Security, while nice, is not a formal requirement. Delivery is. As such, security solutions that the network community wants to implement need to align and support the requirements of the community that will serve to adopt - if they see value for themselves.

In assessing organizational practices of network providers, we find three distinct areas of opportunity where a business is likely incentivized to make an investment or adopt a solution:

Enabling Management and Troubleshooting: Providers earn revenue by maintaining reliable and performant networks for their customers. Solutions that reduce exposure to downtime, proactively inform, or simplify management provide strong value in supporting the core functionality of a business.

Provide Organizational Value: Opportunities for organizations to show added customer value, differentiate from competitors, or generate additional revenue help market and sustain business operations, providing opportunity for business growth.

Establish Evolvable Design: Given a continual advancements in technology, investments that demonstrate modular approaches and capabilities offer a strong financial incentive for adoption. In this regard, technologies that serve as building blocks, which can individually be upgraded as business needs

require, help minimize large recurrent business investments.

By focusing on any of these areas in a solution design, we can demonstrate real value to an organization, helping to incentivize adoption of components that security approaches could leverage at a later time.

4.3 Demonstrating an Incentivized Approach

Before we introduce our architecture, we present one example to demonstrate how a solution could offer initial value to a business, serving as an incentive for adoption.

Underlying the business relationship between customers and providers are Service Level Agreements (SLAs), contractual obligations precisely defining what the service provider will guarantee, what the customer can expect, and remedies should the provider not perform. For transit services, these agreements typically cover a range of network metrics, often including latency, packet delivery, and network availability as core service guarantees [31].

Because SLAs underlie the core functionality of a service provider, efforts to ensure network performance are critical to a businesses success. However, these guarantees are balanced against costs, ensuring that a provider is efficiently providing service without over/under-subscribing the network.

To ensure continued levels of performance while accounting for events which may impact a SLA, network administrators often incorporate redundancy into architectures while balancing network loads to achieve a resilient architecture. The problem with this approach is two-fold. First, it creates a reactive network which solely responds to a *network* event, potentially creating an outcome that is *network optimal*, but *SLA detrimental*, e.g. a planned fail-over link is now oversubscribed causing a SLA detrimental effect. Second, this often requires a provider to pay for extra capacity that they maintain, but fail to utilize efficiently.

4.3.1 Leveraging Data Broadly for Better Business Outcomes

In order to prevent an SLA violation event, we need to leverage both network and business data more broadly to make *SLA optimal* network decisions - decisions that primarily support business objectives first and tailor the network response around these requirements. To demonstrate this approach, we consider three scenarios in which a provider could leverage a global database of network information to

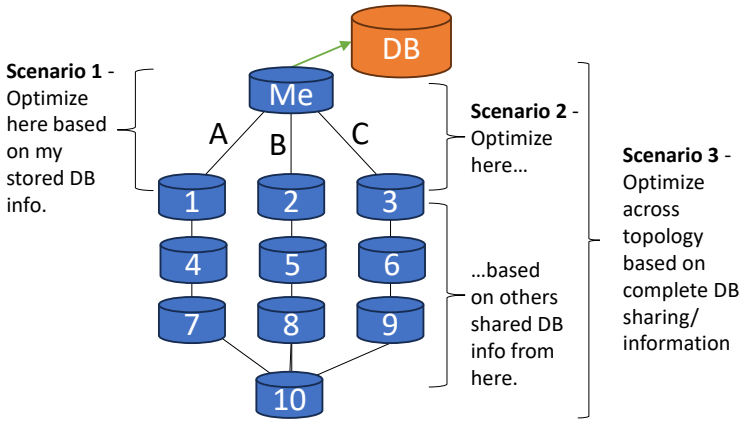

Figure 4.1: **SLA Optimization Based on Network Performance Metrics** With a database of network metrics, we can adjust routing to optimize performance for SLAs dynamically. Value can be realized to a provider where Scenario 1) they are the only participant, Scenario 2) Some level of partial participation occurs, or Scenario 3) a global solution is broadly adopted.

optimize network outcomes while maintaining SLA requirements: 1) the provider is the sole participant utilizing a database of information, 2) the provider participates with a few others in a partial deployment scenario, and 3) full participation in a global database by all parties, as shown in Figure 4.1.

4.3.1.1 Sole Participant

To avoid pitfalls of slow or stagnant deployment, a solution must provide value to the adopter, even if they are the only ones to implement. As a sole participant, a provider would initially utilize the global database to publish and maintain local network statistics, such as live measurements, historical trend data, and other relevant link or route information as needed. In turn, this information can be joined with other business data, such as contract measures of performance, to actively manage and adjust routing behavior at the local level, enabling the provider to maintain optimal forwarding conditions in support of SLA requirements at all times.

Figure 4.1 shows a provider, labeled as "Me", storing historical measurement data of route conditions. This data shows that link "B" is congested in the evening, resulting in customer SLA violations. Normally the provider accounts for this condition through pre-planned or scheduled actions to re-route or balanced excess traffic with link "C". However, link "C" cannot support any traffic due to its own outage (which has been reported to the Database by active measurement and reporting systems). Normally, link

"A" is not used due to high latency along the route that violates a small percentage of SLA agreements. For most SLAs, this path is acceptable. Knowing these link conditions, active network states, and agreement requirements, dynamic configurations which correlate database information can optimize traffic flows across links "A" and "B" without manual intervention or planning, ensuring that no SLA is violated despite ongoing uncertainty in the network.

4.3.1.2 Partial Participation

As participation grows and other ASes begin to share data we can begin to optimize local traffic forwarding based on knowledge external to our AS. For example, "AS 5" reports that the link to "AS 8" has exceeded a 90% capacity threshold and reported an "alert" to the global database. "Me" has subscribed to alerts from "AS 5", being a critical node along its primary forwarding route, and proactively forward traffic to alternate paths, avoiding an SLA violation, and potentially helping "AS 5" until the high usage rate subsides.

4.3.1.3 Full Participation

With broad participation and sharing of data, new paradigms for routing and optimization can be realized. For example, if every node were to share it's transit costs and link performance metrics, we can begin to design routing across a topology that is SLA-optimal at all times while also balancing additional metrics such as cost, capacity, or security. Different thresholds for SLA performance can establish new pricing tiers while ensuring that both obligations and costs are maintained optimally, resulting in topology optimal outcomes in support a differing network objectives and customer requirements.

4.4 Architecture

There are perhaps many forms a final solution could take to incentivize adoption, we present one here - a global database of real-time routing information. This choice is motivated by the realization that underlying RPKI is a database, which administrators of networks have no ability to modify or leverage beyond the intended design. We propose generalizing the database for business use cases first

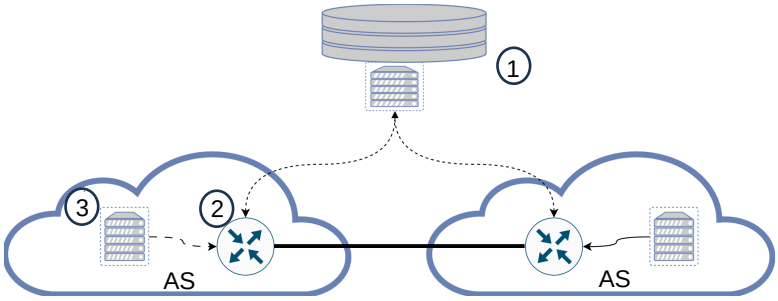

Figure 4.2: **Overall Global Database Architecture.**

to incentivize deployment before establishing security solutions into the architecture.

4.4.1 Overview

Shown in Figure 4.2 is a high-level view of our architecture. Here, a central routing database (label 1) provides a processing and storage capability, along with a set of APIs for read/write access. Each AS interacts through the available APIs to leverage either inline (label 2) (e.g., to support validation or inform routing decisions - similar to the approach used by RPKI) or offline processing (label 3) (e.g., for reporting, configuration staging, or management support).

What to store: We consider three tiers of information for storage. The first, core to the proposal, is routing information as seen by each AS - a lot of insight about the Internet can be gleaned from this, as demonstrated by academic measurements [153, 6], which can be valuable in supporting local routing policy, security, or strategic investments. Second is application specific information. For example, to support an RPKI-like solution, extra information about prefix ownership would be required. Finally, we see private information also being stored. Here, a network provider can tie internal information, like IP address management (IPAM), which can be coupled with routing information to verify route announcements. Of course, this information could be maintained locally by the provider, and present solutions for this already exist [47, 108], but we believe that by broadly decoupling and sharing network data we enable new opportunities for a more secure and dynamic internet.

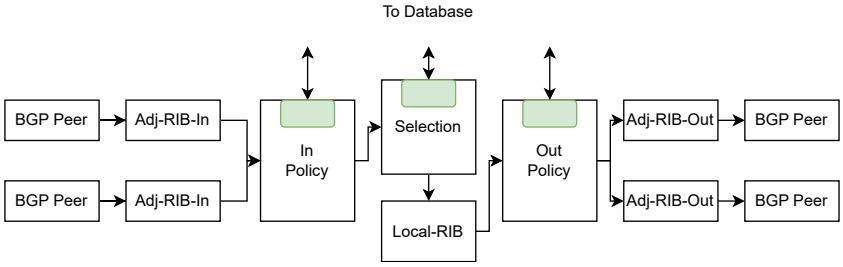

Figure 4.3: **BGP Processing Pipeline.**

4.4.2 Processing

Whereas the database would be a logically centralized entity for information storage, the processing of data is performed at each network provider (and potentially other organizations). We separate this processing into two categories: inline and offline.

Inline processing is logic at the router enabling local routing decisions or providing updates to the global database with the real-time routing information (the core data stored). Here we see two areas in the BGP control plane where extra processing by an AS would be required (illustrated in Figure 4.3 as green boxes). The first are the inbound and outbound policy engine. When a route is received, the inbound policy engine is the initial opportunity to perform an operation. For example, with a route origin validation (ROV) solution, this operation would request a record for the prefix owner from the global database and compare against the recently received announcement before accepting or rejecting the packet. On the outbound side a similar check would be possible. Here, we would compare against the network providers private database of IP address management, verifying no advertisement conflict or leaking of internal routes. Write actions to the global database would also occur at these processing points - e.g. a route advertisement by the AS would update the global database with new routing information, ensuring mechanisms to ease management while ensuring freshness of data.

The second processing point is within the BGP route selection process itself. Here, we believe there is opportunity to make decisions based on what is within the global routing database - i.e., insert new steps in the selection process. Here we could define logic to prefer paths based on which ASes participate in a security solution, which have a history of stability, which route through specific countries, etc. All

of this is opportunity for a network provider to define local value and optimize their network to their or their customer's requirements.

While broad use of inline processing along a network path is sure to raise red flags amongst network purists, there are many ways to mitigate or control these risks, such as default or failover actions, local caching, or use of new routing table designs that incorporate new structures for broader data leverage.

Since routers are not likely to support paths for additional processing immediately, one approach to enable this requirement is with a proxy (which we discuss in our prototype).

Offline processing complements inline processing and is where we see value for network management tasks or pre-establishing route configurations in preparation for action to anticipated network events. We see this as locally deployed software which can read from the database to perform some specific analysis - e.g., to help with troubleshooting. We'll note that much of this information may already be collected and used within present commercial solutions. However, our main proposal is to decouple this network data from the local environment, opening the door for broader use and innovation.

4.4.3 Value to Network Provider:

Central to our approach is prioritization of capability and functionality to incentivize adoption by a provider. By utilizing a global database, a number of opportunities exist for a provider to realize value beyond present localized solutions.

Dynamic Action/Troubleshooting: core to the information database is operational network data which can enable both local and global network insights (similar to what might be provided by a tool like ThousandEyes [179]), providing invaluable information for determining the root cause of events or the operational insight needed to pre-plan network actions and/or respond to external events dynamically.

Organization value: Shared network data can enable new business paradigms, such as the ability to identify and select routes that offer a high degree of security due to identified and published data on security mechanism use. Alternatively, new types of data can be introduced into network decision making, such as business or contract requirements, enabling opportunity for new services, agreements, and value for customers.

Evolvable: the first deployment is always a hurdle, but once in place, this architecture can enable continued evolution by the community without requiring further broad or disruptive change. In fact, most of this infrastructure is already available in various form through RPKI and other network measurment initiatives. However, by centralizing and joining the outputs of these solutions we enable broader value and opportunity without increasing the complexity or overhead incurred through multiple narrower approaches - key considerations for adoption.

4.5 Experimental Approach

To demonstrate both the feasibility and opportunity resulting from a centralized solution, we simulated a network topology of 53 ASNs, each with a local processing capability that could leverage and interact with our global database for a variety of network opportunities or decisions. We first cover the design of the core components necessary to enable our design before highlighting our approach to demonstrate initial value, security, and the requirements necessary to scale to a full deployment.

4.5.1 Setup

4.5.1.1 Global Database Design

For our global database design we chose to deploy MongoDB. MongoDB uses a document-oriented design to organize data, which allowed for individual customization of records as needed for each AS. Horizontal scaling and redundancy would further allow us to adjust our model to larger processing scales without significant redesign of the database itself due to this document-oriented approach [117].

We structured our data documents into a public and private record for each AS, as shown in Figure 4.4. The public record represents information that an AS would be willing to share about their state and relationships, and could be read by any AS in the topology. The private record contains information an AS would want to store and leverage, but not share publicly (e.g. potential private links, relationships, non-public network performance, etc.) Finally, we established accounts for each AS and assigned appropriate permissions to both manage and interact with the various documents by leveraging

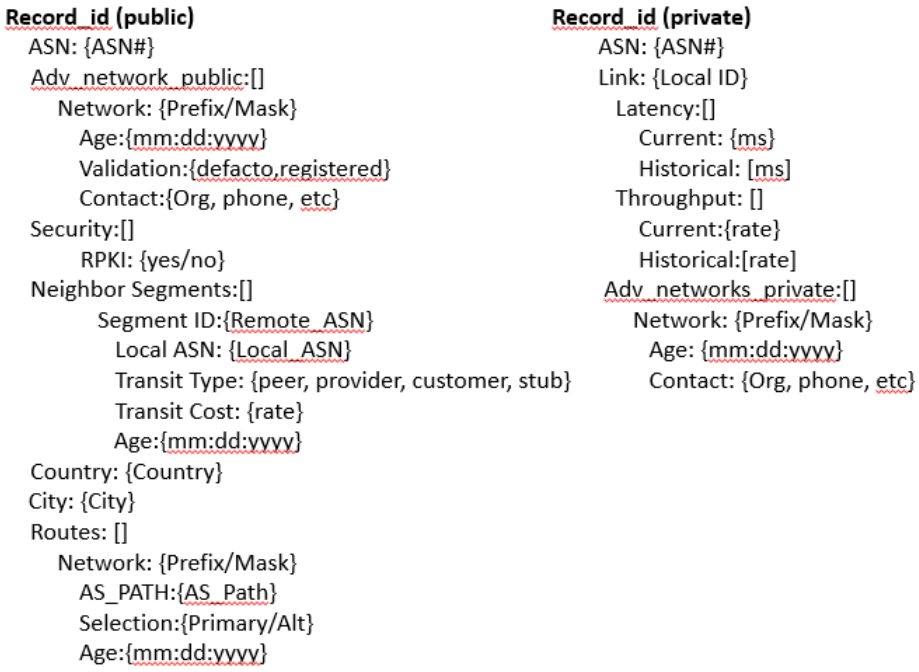

Figure 4.4: **Database Templated Document Data Hierarchy** Both a public and private record is maintained for each AS allowing for both broad public use and select private control.

Mongo's available API's and search structures.

4.5.1.2 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 router to intercept/process packets prior to reaching/leaving the router. For future iterations, we plan to fully incorporate this design into BIRD itself.

To operate, the proxy would monitor for incoming and outgoing packets and conduct processing actions aligned to our stated goals. As an example, if we wanted to only accept routes where 50% or more of the ASNs along that route were deploying RPKI, the proxy would listen for BGP updates containing network layer reachability information (NLRI). When an update containing NLRI information was received, it would then identify the ASes contained in the AS_Path before querying each AS's database record to verify if RPKI was deployed. A final RPKI participation score was then calculated for the

route. If the route did not meet a pre-defined threshold, it would be stripped from the BGP packet with the remainder forwarded to the router for processing. With this approach, we could control which routesto accept or reject based on our defined requirements.

Because modifying BGP packets in-line would adjust TCP sequence metrics, causing de-synchronization of a session between two paired ASes, we implemented local session handler to track and maintain packet sequencing. This allowed us to maintain TCP connections with proper packet handling between nodes regardless of modification. Once implemented in BIRD, we will not need to maintain this feature.

4.5.1.3 Simulation Design

To measure and validate the interaction between the proxies and database, we simulated a network topology of 53 ASNs utilizing the SEED Internet Emulator [154]. SEED leverages Docker containers running the BIRD routing daemon to replicate autonomous systems, allowing for a scalable and flexible network simulation platform. To model our network as closely to real-world as possible, we utilized a subset of the CAIDA AS-relationship dataset to define linkages and peering relationships amongst each AS [53]. Our final environment consisted of a central clique with six primary ASes, six Tier-1 ASes, eighteen Tier-2 ASes, and 23 customer-stub networks dispersed within the topology. Interconnecting these ASes were 232 peer-to-peer and 76 provider-to-customer links which established BGP route forwarding characteristics according to [21]. The resulting average and maximum topology path lengths were 2.6 and 5, providing a reasonable approximation to the real-world average of 5.3 [77], depending on vantage.

Latency was further incorporated by adjusting the queuing discipline for select nodes relevant to our measurement goals. From this topology design our final environment relied on 196 docker containers, as shown in Figure 4.5.

4.5.2 Experimental Approach

We organize our experiments into three overarching themes of value, equivalency, and performance. For value, we show how a global database could be leveraged to provide value to an administrator or their business, helping to incentivize adoption. We then demonstrate how security approaches could

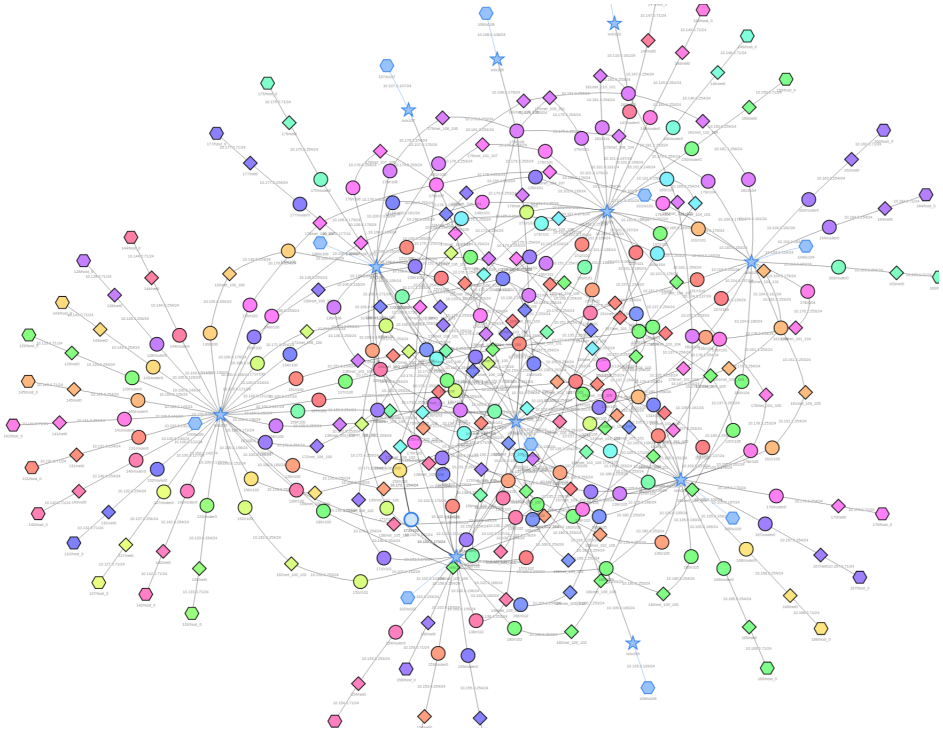

Figure 4.5: **Final Simulated Topology of CAIDA AS-Relationship Dataset Sampling.**

easily extended from this design, providing equivalent capability to present approaches like RPKI and BGPsec while also enabling opportunities for new paradigms. From our results, we conclude with design requirements to effectively achieve our approach at scale.

4.5.2.1 Demonstrating Value

We first align our solution to offer value to a business, a key incentive necessary to encourage adoption. We demonstrate two approaches where a business would likely find initial value: SLA optimization of cost and route selection based on reported security deployments.

SLA Cost Optimized Route Selection. To demonstrate how a provider could leverage a SLA cost optimization approach, our database is initialized with measured and/or publicly available transit costs, demonstrating how data could be included to add value without broad participation. As participation grows, the measured data is replaced with validated source costs (reported directly by the record owner/system), providing a path for stronger assurances over time.

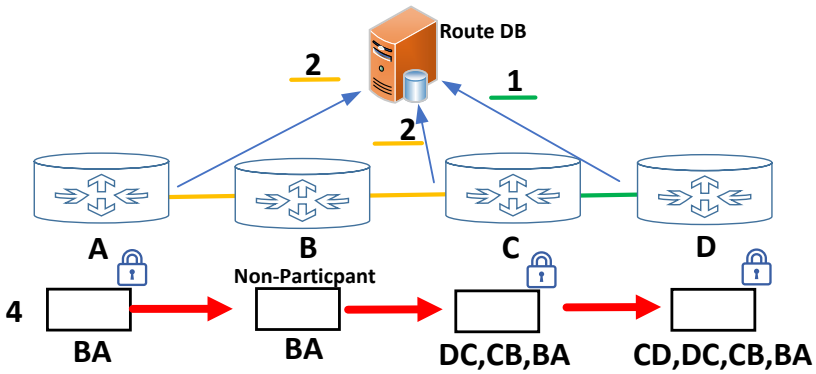

Figure 4.6: **Establishing Path and Topology Validation**

To optimize costs, each router listens for incoming route announcements, calculating an overall transit cost for each route observed based on available data for each AS in the AS_Path of a BGP update. If a route demonstrated a lower cost than the route present in the BIRD BGP table, the proxy would forward the new lower cost route to the router after assigning a pre-defined BGP community value to the received BGP update. This community value would then be used by the router to prioritize the lower-cost route for selection.

To measure this approach in our simulation, we initialized all non-peer-to-peer links according to our chosen topology deployment rate (e.g. 10% represented 10% of peer-to-customer links reporting transit costs) using publicly reported transit costs for 10 Gb/s links in North America [167]. Since actual transit costs can vary based on overall throughput of a link, we assumed a constant 85% usage rate. Future design iterations could account for dynamic price changes relating to traffic rates, if needed.

To measure our results, a dump of the routing table is first conducted and saved for later comparison. Then, each router in the topology was reset (over a 20 minute time period), resulting in BGP updates being generated and received by the proxy, processed, and added to the routing table as appropriate. At the end of the experiment, we would conduct a dump of the routing table, calculate transit costs for each route and compare to the original state to determine the occurrence of lower cost routes being selected.

SLA Compatible Cost-optimal Route Selection Results. Relying on the observed proxy route data across the varying deployment rates (Figure 4.7) we find a number of opportunities where a provider would find value from this approach. First, compared to the default route selected by BGP, up to 15% of

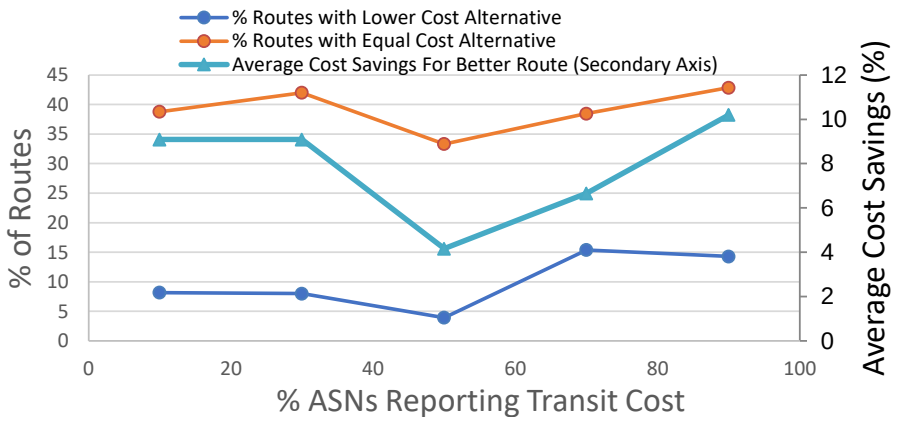

Figure 4.7: **Availability of Cost Alternative SLA-Compatible Routes and Cost Savings Compared to Default BGP Route Selection.**

routes observed to a destination were lower cost. Alternatively, 40% of non-selected routes were equal cost to the chosen route, demonstrating opportunity where a provider could dynamically adjust forwarding without unexpectedly increasing operational costs. On average, achieving these lower cost outcomes required just one additional hop (Figure 4.8), representing a negligible increase in additional latency, an important consideration for SLA agreements.

While a provider would only be able to effect local forwarding decisions, broad participation would open up the opportunity for holistic topology optimizations. However, even if a provider were the only one to use this approach, the ability to incorporate and adjust network actions in support of business objectives at the local level still provides a valuable opportunity for managing business requirements based on metrics of interest to them. Further, while upstream costs are not necessarily directly passed to downstream customers, they are likely to have indirect effect, providing a metric by which to plan with.

Validated Path Route Selection. In order to establish the conditions where we could reference the existence of a path between two ASNs, upon initialization, each ASN would publish their BGP neighbor relationship to the database, as shown in Figure 4.6. Here, we defined two levels of validation based on either one or both ASes attesting to the link between them as follows: First each participating AS would published their neighbor relationships to the database, along with a permission set that allowed the distant neighbor to sign the same record as an attestation. This initial publication represents a partially validated link (label 2). The distant neighbor can then attest to this record by signing the local

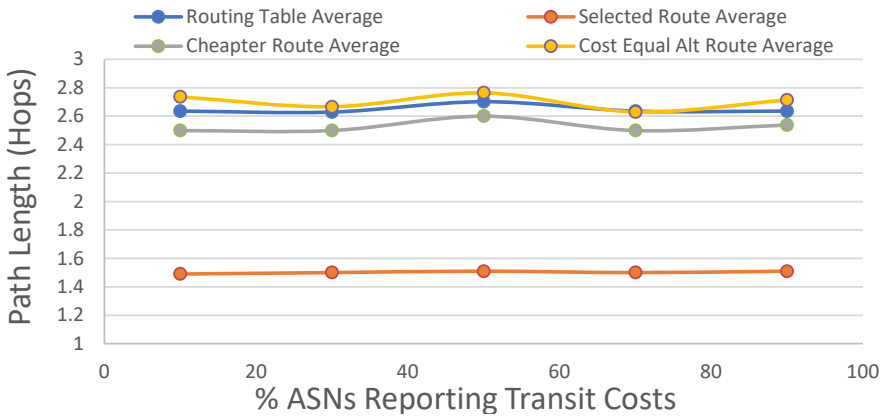

Figure 4.8: **Average Increase in Additional Hops to use a Lower Cost SLA-compatible Route**

AS's path statement utilizing the permissions granted. This path is then be 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 existing, one for each ASN's public record.

Utilizing the same approach presented in the SLA example, each proxy would instead select a route based on the percentage of a path validated, demonstrating an approach where a route could be selected based on a defined metric of security that the provider was interested in.

Validated Path Route Selection Results. To assess value, we again compared the availability of validated routes to those selected by the normal BGP process over a range of participation rates, as shown in (Figure 4.10). At the lowest rates of participation, both lower and equal cost routes were available up to 30% of the time, diverging in favor of equal cost alternatives at higher participation rates. This is mostly a result of the BGP selection process favoring the shortest path to a destination. Alternative routes would typically incur additional hops to select a more secure alternative, making it less likely to receive a route with greater overall security. This is also demonstrated in Figure 4.11 where the overall increase in the percentage of the path being validated was minimal compared to the selected route in the BGP table (less than 10% improvement in percentage of path validated). The high validation in this result was primarily due to the selection of our measurement node being very central in our topology along with the random assignment of the published data for the path occurring on the primary path for this ASN's traffic.

The overall cost to select a route with better validation averaged between 1-1.5 additional hops. How-

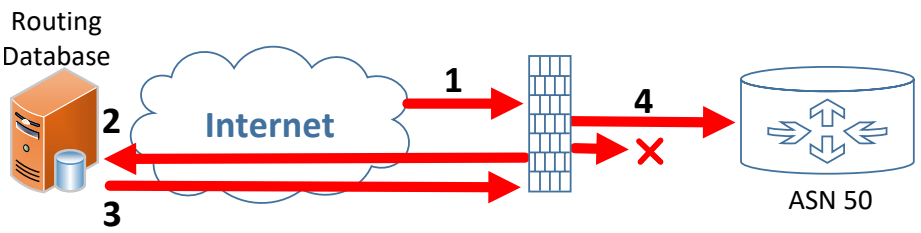

Figure 4.9: **ROV Utilizing Routing Database**

ever, if a provider is able to show that they can offer better opportunity here for a customer, the value of presenting such an approach could incentivize a business by offering opportunities for new customer solutions.

As an alternate approach, we further assessed the availability of routes where a single non-participant was separated two participating nodes, allowing for a "degree of assurance" metric by which to validate a path against, shown in Figure 4.6 as node "B". By using this approach, we could offer a degree of certainty that the path provided existed and that the node between the two participants exists. Use of network time data could further be incorporated to show that no malicious routing actions could have occurred between these two points demonstrating a flexible design to tailor solutions based on local value.

Demonstrating Equivalency.

To demonstrate that our database could perform while providing an equivalent approach to RPKI, we established a comparable approach using our database, as shown in Figure 4.9 (label 1). Here the proxy would act inline, listening for inbound advertisements. For each received advertisement, the proxy would conduct 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) against the data record in order to determine whether or not to accept the announcement (label 4). We then randomly assigned our equivalent approach to a defined percentage of ASes within our topology before executing a prefix hijacking attack. To ensure a direct comparison, we recorded the ASes selected for deployment to ensure the same systems would deploy RPKI on subsequent tests. After the attack, we assessed each router's routing table for presence of the hijacked route.

Demonstrating Equivalency Results. In our comparison we were able to show that utilizing a database could provide equivalency to RPKI while offering broader opportunity to evolve. Across

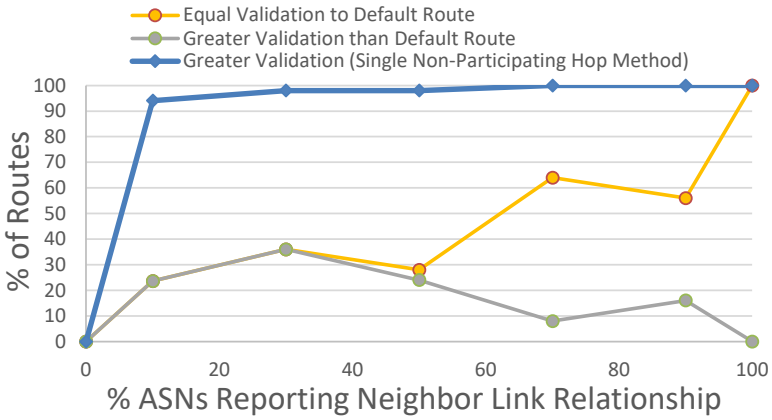

Figure 4.10: **Route Selection Opportunity based on % of Path Validated.** Using random selection for database participation and publication of neighbor links, we demonstrate the opportunity available to select alternative routes based on the percentage of a path validated compared to the default BGP selected path. Availability of fully validated routes using the single gap method is further demonstrated.

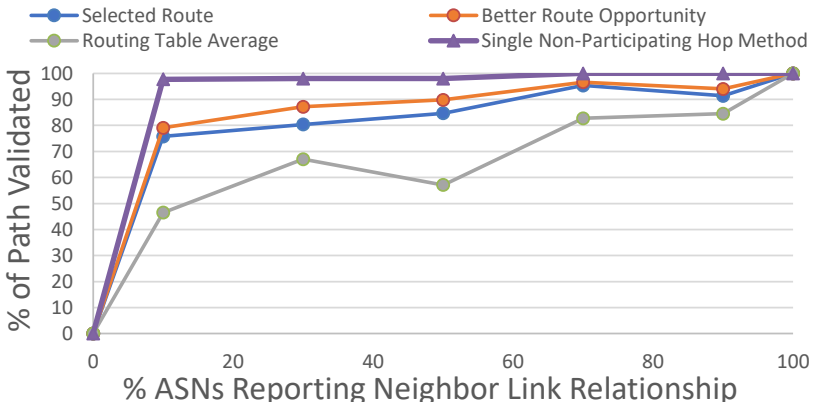

Figure 4.11: **Comparison of Percentage of Path Validated.** Using the same random selection previously, we measured the best route available from our measurement node based on percentage of path validated for different selection methods.

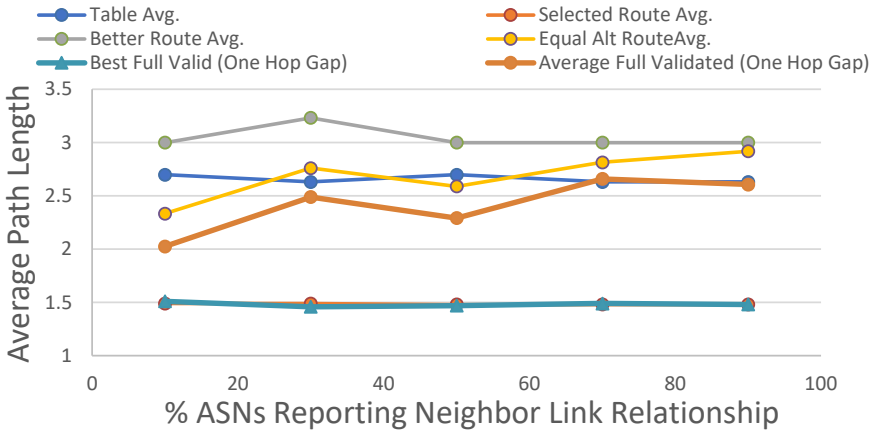

Figure 4.12: **Average Hop Increased to Select a Path with Greater Validation.** Using alternate path selection methods, we demonstrate the average path length compared to the default BGP selected route.

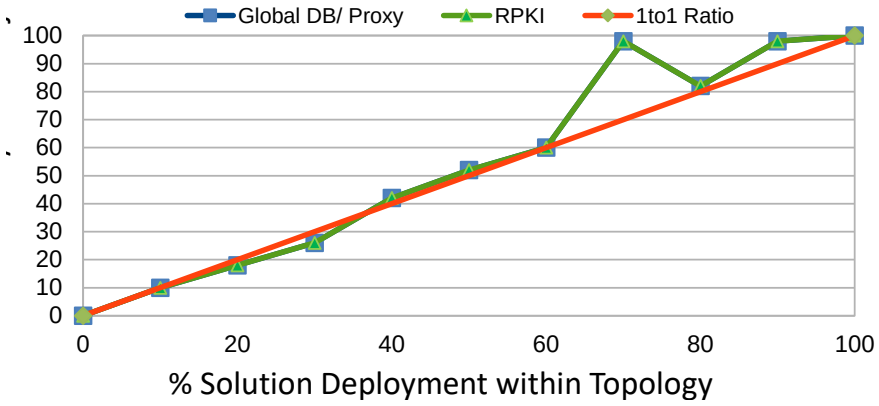

Figure 4.13: **RPKI/Database Approach Comparison.** In a 1to1 comparison, our global DB approach was able to match RPKI results (hidden behind RPKI line) while offering greater flexibility to leverage data for other needs.

the varying deployment scales, we found complete agreement between both approaches, as shown in Figure 4.13. Measurements below the one-to-one line reflect a random selection rounding reduction and not a drop in security, while measurements above reflect security being extended to non-participants in the topology due to an upstream provider rejecting the bad route preventing propagation.

Measuring Topology Performance. In order to understand how our approach would scale to a global deployment, we utilized two approaches for topology performance measurements. In the first approach, we established a single measurement node to gather and record inbound BGP updates, which were generated by restarting the BGP process for each router in the topology. After resetting a router, we allowed time for the resulting updates to traverse and settle before triggering the next event. This allowed us to gather metrics surrounding the amount of requests to the database being triggered for every event along with additional metrics like average path length and prefix counts per update, which we compare against global metrics measured through RIS [145].

Since this approach would only generate limited traffic rates before negatively effecting topology routing, our second approach a traffic generator we built to create BGP updates for higher rates of traffic. This generated update approach measured the overall performance and capacity of the proxy to process BGP updates across varying BGP update rates, which we compared to global rates to understand how our environment compared.

Topology Performance Results. For our topology generated BGP update approach, we obtained

Table 4.1: **Select Topology and BGP Update Metrics.** Topology Metrics captured from simulating an outage event at each router and capturing resulting BGP update propagation.

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

metrics resulting from a complete measurement run, as shown in Figure 4.1. Overall, an average of 4.63 database lookups per BGP update occured, which is slightly lower than our calculated global rate of 6.04 from [45], assuming both a path and prefix validation solution were implemented. Utilizing these averages, we calculated the expected rate of server requests across varying update rates and compared to our results in Figure 4.14. We note two items of interest, first, 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 were possible, up to approximately eight updates/s before topology resets had a significant effect on overall delivery.

To calculate peak throughput for the database, we generated server request packets at a peak rate of 20k requests/s, with a steady state operation averaging approximately 11k requests/s, as shown in Figure 4.14 (secondary axis).

Using our update traffic generator, we measured overall processing time and throughput capacity across varying network latencies, as shown in Figure 4.15. 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, which is the default size offered in most systems, to include our proxy. Results of the maximum throughput and queuing capacity are shown in Figure 4.16. Of note is the significant effect that network latency imposes on overall processing rates, demonstrating that careful consideration of database design and location(s) is(are) necessary. We discuss opportunities to address this in Section 4.6.

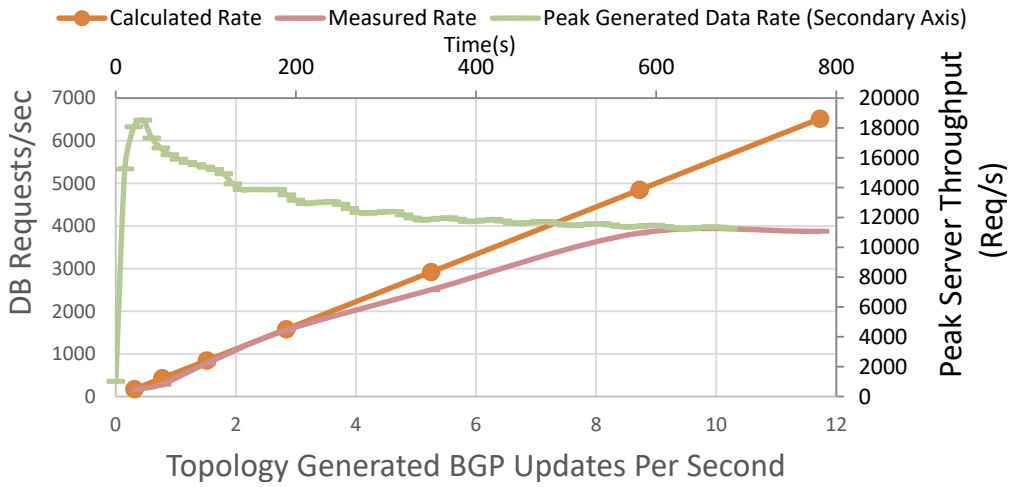

Figure 4.14: **Peak Server Load.** Utilizing topology generated events, server load was measured across various BGP Update rates and compared to calculated rates from packet metrics. Variance between calculated (orage) and measured (pink) are a result of topology generated events having effects on forwarding, resulting in loss of BGP updates. Peak server throughput was measured via a traffic generator to simulate rates beyond what the topology could generate alone.

4.5.2.2 Extrapolating to Global Requirements

For our calculations, we provide two sets of values, the first assesses requirements based on a single request or action 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. From these two approaches we can extrapolate requirements to other designs. For example, to verify the average performance along a route, we would need to request metrics along every node. If more than one metric were desired per node, we would simply double the requirement. For calculations, we rely on the values shown in Table 4.2 and our previously calculated metrics for an average BGP update. Reported metrics for the average number of BGP updates were consistently higher than the globally average. For worst case caluclations of peak loads, we rely on AS65000 reported values. Results of these calculations are shown in Table 4.3.

While it would be nearly impossible to precisely define overall requirements for a global database, owing to variances in adoption, provider use cases, growth of integrated features over time, changes in supporting technologies, and differences in potential architectural approaches to realizing such a solution, the above metrics do present, in our opinion, a viable opportunity for such a design. Methods to lower

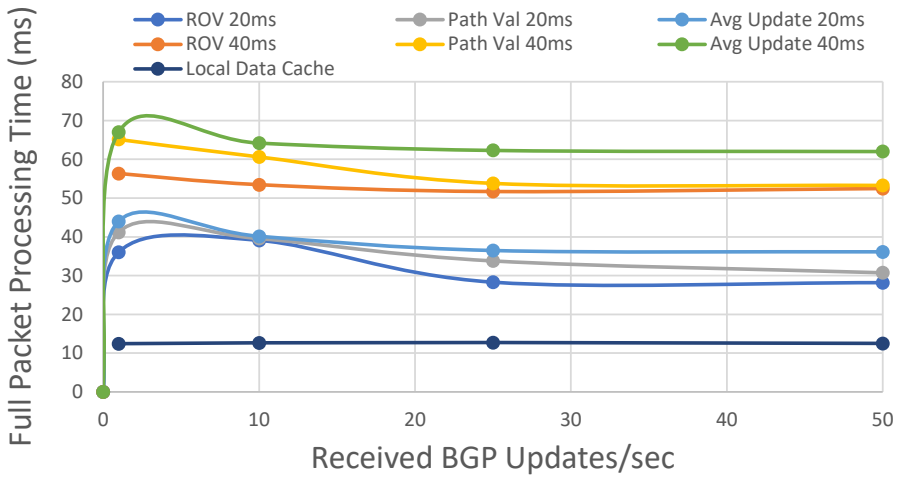

Figure 4.15: **Latency Effect on Processing Time.** Latency Effect on processing time for a packet requiring a single validation event (ROV) and multiple validations (Full Path) are demonstrated. An Average BGP update containing one prefix and an AS-Path length of five is demonstrated and used to calculate overall throughput.

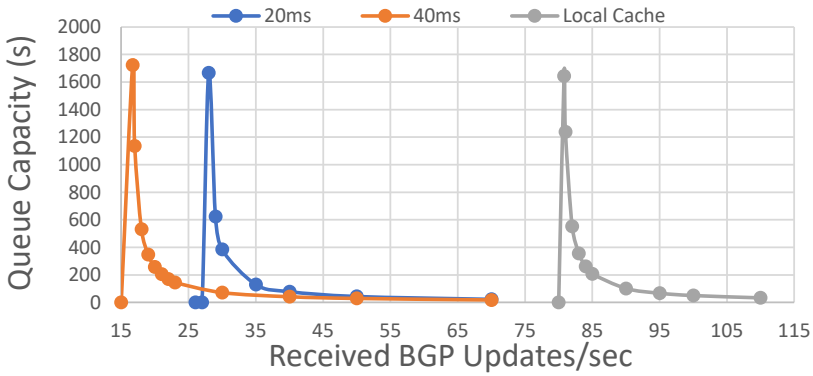

Figure 4.16: **Overall BGP Update Processing Throughput and Packet Queue Capacity.** Peak throughput of 16.12, 27.4, and 80.19 BGP updates/s are demonstrated across select latency's. Queue capacity is based on a maximum queue of 1000 packets, demonstrating potential ability to scale for peak events.

the design requirements could easily be integrated, which we highlight in Section 4.6.

4.6 Discussion

Oour primary focus is centered around offering value to a provider in order to incentivize adoption of solutions that we can later build security approaches off of. We note that this could take on many forms, of which we present only one. At initial glance this may seem like an insurmountable challenge, but modifications such as regionally aligned databases can serve to lower overall requirements. Alternatively,

Table 4.2: **Global BGPUpdate Stats and Single AS Sampling.**
Note: * ASN65000 Values Represent Actual Table Updates for 1hr reporting window and not rate of BGPUpdates received. ** represents a significant one-time anomaly during reporting period. For these events, the average without these anomalies follows in parentheses.

BGP Global Stats [77]		AS65000 Stats [45]
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

Table 4.3: **Calculated Global Database Network Throughput Rate Requirements.**
Note: ** assumes that a single system could perform 20k requests to a database at once. ***Utilizes BGP Update Metrics measured from AS131072 over 14-day period

Server Throughput Requirements for Global Events		
Origin Validation Calculation	**Req. Load**	**Resp. Load**
# ASNs * Event Rate * Packet Size * Avg. Prefix/Update = Server Load		
Average, Even Distribution over 1Hr Sample	1.547 Mb/s	1.255 Mb/s
Avg. Max, Even Distribution over 1Hr Sample	13.76 Mb/s	11.16 Mb/s
Avg. Max, 1sec Distribution (Theoretical)	49.53 Gb/s	40.16 Gb/s
Avg. Max, 1sec Dist. w/ 8.8sec Prop. Delay [56]	5.63 Gb/s	4.56 Gb/s
Peak Anomaly (20k), 1sec Dist. (Theoretical**)	412.7 Gb/s	334.7 Gb/s
Path Validation Calculations	**Req. Load**	**Resp. Load**
# ASNs*Event Rate*Packet Size*Path Length* Avg. Path/Update = Server Load		
Average Rate, Even Dist./1 Hr, Avg Len.	7.7 Mb/s	6.2 Mb/s
Average Rate, Even Dist./1 Hr, Max Len.	18.9 Mb/s	15.34 Mb/s
Avg. Max, 1sec Distribution, Avg. Len.	128.1 Gb/s	103.9 Gb/s
Avg. Max, 1sec Distribution, Max Len.	313.3 Gb/s	254 Gb/s
Avg. Max, 1sec Dist. w/ 8.8sec Prop., Max Len.	35.6 Gb/s	28.9 Gb/s
Peak Anomaly (42.5k), 1sec Dist., Max Len.	9.65 Tb/s	7.82 Tb/s
Average BGP Update	**Req. Load**	**Resp. Load**
# ASNs*Rate*Size*Path Length* Avg. Update Metric*** = Server Load		
Average BGP Update Rate, Path & Prefix Count	688 Mb/s	558 Mb/s
Peak BGP Update Rate, Avg. Path & Prefix	1.81 Tb/s	1.47 Tb/s
Average Prefix Update Rate & Path	22.5 Gb/s	18.8 Gb/s
Peak Prefix Update Rate & Path	434.3 Gb/s	363.2 Gb/s

if the RPKI infrastructure could be extended for broader use, it could potentially provide an avenue to build off-of while also leveraging present adoption efforts. Regardless, to have a successful effort, a solution needs to be centered around offering providers immediate value or solutions to challenges they face as incentives for adoptions. While measurement infrastructure, network monitoring, and security approaches exist and are used by providers already, the ability to correlate this data effectively across disparate systems is a challenge. That is where we see the value of a centralized approach offering value. Further, once the initial deployment hurdle is over, we enable a platform that can easily evolve to new

ideas or approaches without significant recurring investment - an incentive in and of itself.

4.7 Future Work

A number of design improvements are considered for future implementation. While we utilized an inline proxy for a proof of concept, a viable design would integrate with the routing system itself to leverage a number of efficiencies. We plan to incorporate our approaches into BIRD, allowing us to easily reference local tables, conducting lookups to a database only when the result would effect a change in the local state. This can serve to limit the overall requests a global database would have to serve. Second, with an integrated design, we can leverage local actions that an administrator would normally perform, such as adding a new BGP neighbor, to automatically trigger validation checks as a preventative measure. Further, this can also serve to trigger database updates, ensuring freshness without requiring additional involvement by an administrator, reducing management overhead. Many other thoughts are also considered.

4.8 Related Works

While our work is uniquely focused on establishing network provider value as a direction to incentivize adoption, a number of research directions are closely related to our underlying structure.

Internet measurement methods have proposed a number of approaches to gathering broad network data in support of research, visualization, and troubleshooting [156, 48, 171]. However, these approaches work to reverse engineer networks, inferring architectures from limited vantages. A stronger approach would be to focus on establishing the infrastructure necessary to enable broad community participation in reporting, sharing, and leveraging network data more broadly, which our work uniquely argues for. However, these measurements are still important, and could serve as initial data within this architecture, helping to establish initial value but with broader opportunity for use.

Security solutions, particularly with BGP, identify a number of novel approaches and methods which are best captured in a number of survey papers [116, 177]. A common challenge amongst these efforts is a narrowly focused solution that solves only a portion of the challenges present with BGP, resulting in limited adoption.

Recently, use of blockchains as a database structure to establish mechanisms for BGP security have been proposed. Here, approaches for tamper resistant BGP messaging [69] and resource management [195, 72, 135], with functionality similar to RPKI, have been demonstrated. One problem with these approaches has been the throughput necessary to achieve internet scale requirements. Authors of Routechain [150]proposed a bi-hierarchical blockchain structure to alleviate some of these performance considerations, but this remains an active area of research.

4.9 Conclusion

Our proposal is simple, we should focus on building solutions to better enable network administrators to perform their core functions, as they are the ones who will ultimately deploy new network solutions. By focusing solutions on the administrator, we can establish core components for easier deployment of security mechanisms, leading to broader adoption. Further, we must make network data more broadly available to leverage new routing paradigms or business opportunities. We are artificially constraining opportunity by proposing solutions that are either narrowly scoped or static in their design. Rather, we need to approach future network development from a perspective that will allow broad opportunity to evolve and shift with changing perspectives, timelines, and directions. We demonstrated one such approach, a global routing database, which enables opportunity for an evolvable design. While open questions remain, we hope our demonstration identifies new research directions for the community by which to enable a stronger security environment for the Internet.

Chapter 5

Discussion

Incentivization is not a singular approach to encouraging security adoption. Rather, it is an overarching method that must be tailored to each environment commensurate with the overall goals. For example, the home router environment could be incentivized to use a security package that is publicly available rather than creating their own - if the right caveats for commercial use exist. This would help them quickly incorporate features into their system without having to conduct significant design of their own. However, this introduces second or their layers of impact that need consideration such as package management, code validation, etc. If the secondary costs become too much, the initial incentive may no longer provide the envisioned value.

Alternatively, a manufacturer could be incentivized to secure their system through the application of cost, e.g. where not complying imposes a cost that would help incentivize adoption. Regulations are one such approach. Here non-compliance can prevent an organization from selling a product or impose a fine. At present, there is little regulation governing home router security in the U.S. A select number of state laws, such as California's law mandating default password changes as part of setup, exist, but their scope is narrow compared to the breadth of exposures found in systems. A larger scope regulation could, in our opinion, go a long way to addressing exposures in the home gateway segment as a strong incentive for compliance. In this regard, the european market has shown strong approaches to governing the digital landscape with privacy protecting regulations, such as with GDPR. These efforts have been successful in incentivizing change across the data environment, showing that regulations do have a strong role in incentivizing compliance.

The best way to incentivize outcomes, however, is by creating value such that people or companies

will want to adopt on their own. To do so effectively, a solution should be designed such that an administrator receives value, even if they are the only ones to implement. While we demonstrate one approach in this dissertation using a global database, any approach which offers clear value could be equally valid. The core point of this method, however, is that we need to stop approaching solutions as standalone efforts that incur costs without also considering the incentive to offer value and encourage adoption. This is where present approaches, such as RPKI and BGPsec, fail, as they impose administrative costs up front while delaying the received benefits until broad adoption has occurred. This mismatched cost/value approach has enabled the present stagnated state of adoption for RPKI while preventing others, such as BGPsec, from ever being implemented.

Chapter 6

Conclusion

The challenge of broad security adoption does not have to remain a distant goal. With the right use of incentivization, through the alignment of value, we can better encourage adoption of solutions, either directly, such as with regulation, or indirectly, as demonstrated through our approach for a global database. While there are certainly many other approaches that could follow differing structures, our core point is that we are incorrectly focused on building security solutions that solve a security issue but are not aligned with the goals of the demographic needed to adopt. In many cases, these approaches impose up front costs, either technical or managerial, while delaying the value until broad adoption is achieved. While altruistic goals for internet security are admirable, it is unlikely that we will ever see complete adoption of a solution using this approach. The stagnation of RPKI only serves to bolster this position.

As we demonstrate in our final effort, building incentivization into a solution helps to change this value proposition such that the demographic looking to implement would likely be encouraged to do so. By building value into a solution, or by focusing on problems that administrators care about, we help to align the goals of both parties (administrators and the security community) while making better outcomes for both. Uniquely, this approach allows the security community to then leverage the opportunities created, as part of a planned design incorporation, approaches for security that do not impose additional or new costs and/or administrative overhead.

Appendix A

Research Efforts

A.1 Primary Publications

(1) **Title:** Natting Else Matters: Evaluating IPv6 Access Control Policies in Residential Networks

Conference: 22nd International Conference on Passive and Active Measurement 2021.

Research Effort: This project assessed consumer gateways in order to understand the security approaches used by gateway manufacturers for IPv6 access controls. Our work identified significant security shortcomings across manufacturers and devices. Further, we reported three significant gaps in security to these manufacturers, of which two were patched and one was never fixed. Overall, we find an environment of poor support for IPv6 with access control mechanisms which often fail to account for operational differences compared IPv4, such as ephemeral addressing. In a large percentage of cases, consumers networks were left exposed out-of-the-box with no adequate control mechanism available to remedy.

(2) **Title:** Doomed to Repeat? A Historical Survey of Consumer Gateway Access Control Failures Under NAT with Correlations to Present IPv6 Deployments.

Journal: ACM Computing Surveys (2023).

Research Effort: This project undertook a historical review and assessment on the use, challenges, and security of network address translation use in home gateways. From this, we are able to correlate lessons to present IPv6 design approaches, showing a similar patter of mistakes occurring with IPv6 deployment in regard to standardization, security, and implementation. Unique to this paper is the historical assessment and correlation of vulnerabilities over a lifetime

of use with regards to NAT in the home gateway environment.

(3) **Title:** DnD-db: A Democratized Network Data Database for Tailored Routing and Security Campaigns.

Conference: Under Submission

Research Effort: To better incentivize security adoption, solutions must incorporate incentives into design in order to encourage administrator implementation. Present approaches, such as RPKI, add costs for administrators while delaying the intended benefit until broad adoption by the internet community is realized - a significant challenge for an environment defined by dis-aggregated participation. For this we present a network data database which can be leveraged to tailor local value leveraging broad availability of network data for local decision making. From this approach, we further demonstrate how security can extend from such a solution.

A.2 Poster/Abstracts:

(1) **Title:** All Your Data Are Available to Us: A Need for Network Segmentation with IoT Devices.

Conference: Networked and Distributed System Security Symposium (NDSS) 2020.

Research Effort: This project investigated segmentation within home networks and IoT devices. While protocols like Universal Plug and Play provide for automatic configuration for hole punching in a gateway, they do not provide capability to segment network access within the local network. This presents broad network exposure within the home where everything operates within a shared trust environment. This project proposed building into the UPnP protocol methods to automate the segmentation of IoT devices within the home and prevent this exposure.

Authors: Olson, K. and Scaife, N.

(2) **Title:** Federating Trust: Network Orchestration for Cross-Boundary Zero Trust.

Conference: ACM Special Interest Group on Data Communication (SIGCOMM) 2021.

Research Effort: This project proposed an approach to federate Zero-Trust decision properties across disparate network standards. By utilizing a proxy handler, networks could coordinate access standards for users without revealing information about the local security environment

to the external entity. This approach would allow companies to maintain threshold security standards necessary to meet security goals of Zero-trust across boundaries, even if they did not have any say in the remote networks security standards.

Authors: Olson, K. and Keller, E.

A.3 Workshops

(1) **Title:** Infinity: A Scalable Infrastructure for In-Network Applications

Conference: IEEE International Symposium on Integrated Network Management 2021.

Research Effort: In-network compute resources are commonly fixed, resulting in a limited or inefficient use of resources for things like edge computing. While computers typically have abstractions that can scale resources efficiently to meet application needs, these same abstraction do not currently reside in network compute capabilities. This paper proposes an approach to scale in-network resources either horizontally or vertically in order to meet the growing compute demands of in-network processing.

Authors: Abranches, M., Olson, K. Keller, E.

(2) **Title:** Enabling Security Research through Efficient Partial Deployment Topology Configuration and Validation

Conference: IEEE INFOCOM Computer and Networking Experimental Research using Testbeds (CNERT) Workshop 2023.

Research Effort: Measuring partial deployment value of a security protocol at internet scale is an impossible problem due to uncertainty in network design and challenges establishing the necessary resources to fully simulate. We present an approach to measure partial deployments quickly and easily and demonstrate our solution utilizing RPKI. From this we show that RPKI requires large adoption rates greater than 70% before the internet will broadly begin to realize the value of RPKI.

Authors: Alharbi, B., Olson, K. Keller, E.

Bibliography

[1] A. Abdi-Nur. Smart TV Upgrade, Privacy Downgrade? In Journal of The Colloquium for Information Systems Security Education, volume 5, pages 22–22, 2017.

[2] B. Aboba and E. Davies. RFC 4924: Reflections on Internet Transparency. BCP 148, April 2007.

[3] B. Aboba and W. Dixon. RFC 3715: IPsec-Network Address Translation (NAT) Compatibility Requirements. Technical report, March 2004.

[4] A. Adams and M. Sasse. Users Are Not the Enemy. Communications of the ACM, 42(12):40–46, 1999.

[5] Akamai. UPnProxy: Blackhat Proxies via NAT Injections. Available from: `https://www.akamai.com/content/dam/site/en/documents/research-paper/upnproxy-blackhat-proxies-via-nat-injections-white-paper.pdf`., 2018. Last Accessed: 17 August 2021.

[6] T. Alfroy, T. Holterbach, and C. Pelsser. MVP: Measuring Internet Routing from the Most Valuable Points. In Proceedings of the 22nd ACM Internet Measurement Conference, pages 770–771, 2022.

[7] B. Alharbi, K. Olson, and E. Keller. Enabling Security Research Through Efficient Partial Deployment Topology Configuration and Validation. In 2023 IEEE Conference on Computer Communications Workshops, pages 1–6. IEEE, 2023.

[8] W. Alliance. WiFi Alliance Wireless Specifications. Available from: `https://www.wi-fi.org/discover-wi-fi/specifications`, 2022. Accessed: 2022-11-17.

[9] Amazon.com. Amazon Sales Popularity - Computer Routers (2020). Available from: `https://web.archive.org/web/20201023233343/https://www.amazon.com/gp/bestsellers/pc/300189/ref=zg_b_bs_300189_1`. Last accessed 23 Oct 2020.

[10] Amazon.com. Amazon EERO Technical Specification. Available from: `https://support.eero.com/hc/en-us/articles/209962973-Frequently-asked-security-questions`, 2021. Last accessed: 2021-02-10.

[11] M. Antonakakis and others. Understanding the Mirai Botnet. In USENIX - 26th Security Symposium, pages 1093–1110, 2017.

[12] C. Aoun and E. Davies. Reasons to Move the Network Address Translator-Protocol Translator (NAT-PT) to Historic Status. Technical report, RFC 4966, 2007.

[13] M. Apostolaki, A. Zohar, and L. Vanbever. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In 2017 IEEE symposium on security and privacy (SP), pages 375–392. IEEE, 2017.

[14] Apple. How to Share Your Wi-Fi Password from Your iPhone, iPad, or iPod Touch. Available from: https://support.apple.com/en-us/HT209368, 2020. Last Accessed: 28 March 2020.

[15] C. G. Assembly. California State Law SB-327. Available from: https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB327, 2018. Last accessed: 13 August 2020.

[16] Audet, F. and Jennings, C. RFC 4787: Network Address Translation (NAT) Behavioral Requirements for Unicast UDP. BCP 127, January 2007.

[17] K. Awe, Y. Malik, P. Zavarsky, and F. Jaafar. Validating BGP Update Using Blockchain-based Infrastructure. Decentralised Internet of Things: A Blockchain Perspective, pages 151–165, 2020.

[18] M. Boucadair, R. Penno, and D. Wing. RFC 6970: Universal Plug and Play (UPnP) Internet Gateway Device - Port Control Protocol Interworking Function (IGD-PCP IWF). RFC 6970, July 2013.

[19] R. Bush and R. Austein. RFC 8210: The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1. RFC 8210, Sept. 2017.

[20] Y. Cadirci. CallStranger Technical Report. Available from: https://github.com/yunuscadirci/CallStranger/blob/master/CallStranger%20-%20Technical%20Report.pdf., 2020. Last Accessed: 17 August 2021.

[21] M. Caesar and J. Rexford. BGP Routing Policies in ISP Networks. IEEE network, 19(6):5–11, 2005.

[22] M. Casado and M. Freedman. Illuminating the Shadows: Opportunistic Network and Web Measurement. Available from: http://illuminati.coralcdn.org/stats, 2006.

[23] G. Chalhoub and A. Martin. But is it exploitable? exploring how router vendors manage and patch security vulnerabilities in consumer-grade routers. In Proceedings of the 2023 European Symposium on Usable Security, pages 277–295, 2023.

[24] E. Chen and Y. Rekhter. RFC 5291: Outbound Route Filtering Capability for BGP-4. Technical report, 2008.

[25] W. Chen, Y. Huang, and H. Chao. NAT Traversing Solutions for SIP Applications. EURASIP Journal on Wireless Communications and Networking, 2008:1–9, 2008.

[26] S. Cheshire and M. Krochmal. RFC 6886: NAT Port Mapping Protocol (NAT-PMP). Tech Report, 2013.

[27] S. Cho, R. Fontugne, K. Cho, A. Dainotti, and P. Gill. BGP Hijacking Classification. In 2019 Network Traffic Measurement and Analysis Conference (TMA), pages 25–32. IEEE, 2019.

[28] T. Chown, J. Ed., Arkko, A. Brandt, O. Troan, and J. Weil. RFC 7368: IPv6 Home Networking Architecture Principles. RFC 7368, Internet Engineering Task Force, October 2014.

[29] CISCO. BGPStream - BGP Outage Monitor. Available from: https://bgpstream.crosswork.cisco.com/, 2020. Last accessed: 13 August 2020.

[30] Cisco. Meraki Go: Easy Networking for Busy People. Available from: https://www.meraki-go.com/, 2020. Last accessed: 18 Oct 2020.

[31] Cogent Communications. Global Network SLA. Available from: www.cogentco.com/files/docs/network/performance/global_sla.pdf, 2021.

[32] F. Communications. Frontier Home Internet Setup Guide (2020). Available from: https://frontier.com/~/media/HelpCenter/Documents/internet/installation-setup/hsi-self-install-guide.ashx?la=en. Last Accessed: 18 Oct 2020.

[33] B. Computer. BGP Outage News Events. Available from: https://www.bleepingcomputer.com/tag/bgp, 2023. Last accessed: 1 September 2023.

[34] A. Constantinescu, V. Croitoru, and D. Cernaianu. NAT/Firewall Traversal for SIP: Issues and Solutions. In International Symposium on Signals, Circuits and Systems, volume 2, pages 521–524. IEEE, 2005.

[35] M. Corporation. CVE-2006-2559. Available from: https://cve.mitre.org, 2006.

[36] M. Corporation. CVE-2012-0383. Available from: https://cve.mitre.org, 2012.

[37] M. Corporation. CVE-2013-3182. Available from: https://cve.mitre.org, 2013.

[38] M. Corporation. CVE-2013-6949. Available from: https://cve.mitre.org, 2013.

[39] M. Corporation. CVE-2017-17746. Available from: https://cve.mitre.org, 2017.

[40] M. Corporation. CVE-2017-7405. Available from: https://cve.mitre.org, 2017.

[41] M. Corporation. CVE-2020-16894. Available from: https://cve.mitre.org, 2020.

[42] M. Corporation. CVE-2020-25988. Available from: https://cve.mitre.org, 2020.

[43] M. Corporation. CVE-2007-2390. Available From: https://cve.mitre.org, 2021.

[44] J. Czyz, M. Luckie, M. Allman, and M. Bailey. Don't Forget to Lock the Back Door! A Characterization of IPv6 Network Security Policy. In Proceedings of the 23rd Annual Network & Distributed System Security Symposium (NDSS '16), San Diego, California, USA, February 2016.

[45] A. P. Data. BGP Instability Report. Available from: https://bgp.potaroo.net/as2.0/bgp-active.html, 2024. Last Accessed: 18 January 2024.

[46] N. De Leon. Many Wireless Routers Lack Basic Security Protections, Consumer Reports' Testing Finds. Available from: https://www.consumerreports.org/wireless-routers/wireless-routers-lack-basic-security-protections/, 2019. Last accessed: 20 Oct 2020.

[47] Domotz. Domotz. Available from: https://www.domotz.com/, 2024. Last accessed: 2024-01-18.

[48] R. Durairajan, S. Ghosh, X. Tang, P. Barford, and B. Eriksson. Internet Atlas: A Geographic Database of the Internet. In Proceedings of the 5th ACM workshop on HotPlanet, pages 15–20, 2013.

[49] M. Dylan. Security Cameras Vulnerable to Hijacking. Available from: https://hacked.camera/, 2018. Last accessed: 18 Oct 2020.

[50] Egevang, K. and Francis, P. The IP Network Address Translator (NAT). RFC 1631, May 1994.

[51] P. Emami-Naeini and others. An Informative Security and Privacy "Nutrition" Label for Internet of Things Devices. IEEE Security & Privacy, 20(2):31–39, 2021.

[52] J. F. and L. Painsignon. RPKI and BGP: our path to securing Internet Routing. 2018. Last accessed: 20 December 2021.

[53] C. for Applied Data Analysis(CAIDA). AS-Relationships Dataset. Available from: `https://www.caida.org/catalog/datasets/as-relationships-geo/`, 2024. Last accessed: 2024-01-18.

[54] B. Ford, P. Srisuresh, and D. Kegel. Peer-to-Peer Communication Across Network Address Translators. In USENIX Annual Technical Conference, General Track, pages 179–192, 2005.

[55] U. Forum. UPnP Specification, 2020. Available from: `https://openconnectivity.org/developer/specifications/upnp-resources/`.

[56] A. Garcia-Martnez and M. Bagnulo. Measuring BGP Route Propagation Times. IEEE Communications Letters, 23(12):2432–2436, 2019.

[57] M. Ghiglieri, M. Volkamer, and K. Renaud. Exploring Consumer Attitudes of Smart TV Related Privacy Risks. In International Conference on Human Aspects of Information Security, Privacy, and Trust, pages 656–674. Springer, 2017.

[58] R. Giobbi. UPnP Enabled By Default, SEI Vulnerability Note VU347812, 2008. https://www.kb.cert.org/vuls/id/347812/.

[59] S. Goldberg, S. Halevi, A. Jaggard, V. Ramachandran, and R. Wright. Rationality and traffic attraction: Incentives for honest path announcements in BGP. In Proceedings of the ACM SIGCOMM 2008 conference on Data communication, pages 267–278, 2008.

[60] Google. Per-Country IPv6 Adoption. Available from: `https://www.google.com/intl/en/ipv6/statistics.html#tab=per-country-ipv6-adoption`, 2020. Last Accessed: 20 March 2021.

[61] Google. Google Nest Technical Specification. 2021.

[62] S. Guha, K. Biswas, et al. RFC 5382: NAT Behavioral Requirements for TCP. RFC 5382, 2008.

[63] S. Guha and P. Francis. Characterization and Measurement of TCP Traversal Through NATs and Firewalls. In Proceedings of the 5th ACM SIGCOMM Conference on Internet Measurement, pages 18–18, 2005.

[64] Guha, S. and Biswas, K. and others. RFC 5508: NAT Behavioral Requirements for ICMP. BCP 148, April 2009.

[65] F. Guo and T. Chiueh. Sequence number-based MAC address spoof detection. In International Workshop on Recent Advances in Intrusion Detection, pages 309–329. Springer, 2005.

[66] T. Hain. RFC 2993: Architectural Implications of NAT. RFC 2993, Internet Engineering Task Force, November 2000.

[67] G. Halkes and J. Pouwelse. UDP NAT and Firewall Puncturing in the Wild. In International Conference on Research in Networking, pages 1–12. Springer, 2011.

[68] N. Hampton and P. Szewczyk. A Survey and Method for Analysing SOHO Router Firmware Currency. 2015.

[69] A. Hari and T. Lakshman. The Internet Blockchain: A Distributed, Tamper-resistant Transaction Framework for the Internet. In Proceedings of the 15th ACM workshop on hot topics in networks, pages 204–210, 2016.

[70] W. Hassan et al. Current Research on Internet of Things (IoT) Security: A Survey. Computer Networks, 148:283–294, 2019.

[71] S. Hatonen, A. Nyrhinen, L. Eggert, S. Strowes, P. Sarolahti, and M. Kojo. An Experimental Study of Home Gateway Characteristics. In Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement, pages 260–266, 2010.

[72] G. He and others. ROAchain: Securing Route Origin Authorization with Blockchain for Inter-domain Routing. IEEE Transactions on Network and Service Management, 18(2):1690–1705, 2020.

[73] R. Heartfield, G. Loukas, S. Budimir, A. Bezemskij, J. Fontaine, A. Filippoupolitis, and E. Roesch. A taxonomy of Cyber-physical Threats and Impact in the Smart Home. Computers & Security, 78:398–428, 2018.

[74] T. Hlavacek, H. Shulman, and M. Waidner. Not All Conflicts Are Created Equal: Automated Error Resolution in RPKI Deployments. In IEEE INFOCOM 2021-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), pages 1–2. IEEE, 2021.

[75] M. Holdrege and P. Srisuresh. RFC 3207: Protocol Complications with the IP Network Address Translator (NAT). 2001.

[76] S. Holzapfel, M. Wander, A. Wacker, and T. Weis. SYNI-TCP Hole Punching Based on SYN Injection. In 10th International Symposium on Network Computing and Applications, pages 241–246. IEEE, 2011.

[77] G. Houston. BGP in 2023. Available from: `https://blog.apnic.net/2024/01/10/bgp-in-2023-bgp-updates/`, 2024. Last accessed: 2024-01-18.

[78] C. Huitema. RFC 4380: Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs). RFC 4380, February 2006.

[79] S. Hunag, Q. Wu, et al. A Survey of NAT Behavior Discovery in VOIP Applications. Journal of Internet Technology, 12(2):199–210, 2011.

[80] International Organization for Standardization. ISO/IEC: 27000:2018 Information Technology Security Techniques Information Security Management Systems Overview and Vocabulary. Technical report, 2018.

[81] S. H.-B. Jackson. California Senate Bill SB-327, Chapter 866. Available from: `https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB327`, 2017. Accessed: 2022-10-17.

[82] C. Jennings. Draft RFC: NAT Classification Test Results. Internet-Draft draft-jennings-behave-test-results-04, Internet Engineering Task Force, July 2007. Work in Progress.

[83] D. Johnson and B. Hartpence. A Re-examination of Network Address Translation Security. 2010.

[84] S. Johnson and others. The impact of IoT Security Labelling on Consumer Product Choice and Willingness to Pay. PloS one, 15(1):e0227800, 2020.

[85] C. Karat, C. Brodie, and J. Karat. Usability Design and Evaluation for Privacy and Security Solutions. Security and Usability, pages 47–74, 2005.

[86] J. Karlin, S. Forrest, and J. Rexford. Pretty good BGP: Improving BGP by Cautiously Adopting Routes. In Proceedings of the 2006 IEEE International Conference on Network Protocols, pages 290–299. IEEE, 2006.

[87] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (S-BGP). IEEE Journal on Selected areas in Communications, 18(4):582–592, 2000.

[88] A. Keranen, C. Holmberg, and J. Rosenberg. RFC 8445: Interactive Connectivity Establishment (ICE). RFC 8445, July 2018.

[89] K. Killourhy, R. Maxion, and K. Tan. A Defense-Centric Taxonomy Based on Attack Manifestations. In International Conference on Dependable Systems and Networks, pages 102–111. IEEE, 2004.

[90] M. Kiravuo, T.and Sarela and J. Manner. A survey of Ethernet LAN security. IEEE Communications Surveys & Tutorials, 15(3):1477–1491, 2013.

[91] B. Kitchenham and P. Brereton. A Systematic Review of Systematic Review Process Research in Software Engineering. Information and Software Technology, 55(12):2049–2075, 2013.

[92] C. Kolias, G. Kambourakis, A. Stavrou, and J. Voas. DDoS in the IoT: Mirai and Other Botnets. 50(7):80–84, 2017.

[93] E. Kovacs. Cybercriminals Steal Cryptocurrency Via BGP Hijacking. Available from: `https://www.securityweek.com/cybercriminals-steal-cryptocurrency-bgp-hijacking/`, 2023. Last accessed: 1 September 2023.

[94] D. Kumar et al. All Things Considered: An Analysis of IoT Devices on Home Networks. In USENIX - 28th Security Symposium , pages 1169–1185, 2019.

[95] C. Kuo, A. Perrig, and J. Walker. Security configuration for non-experts: A case study in wireless network configuration. In Social and Human Elements of Information Security: Emerging Trends and Countermeasures, pages 179–195. IGI Global, 2009.

[96] M. Lepinski and S. Kent. RFC 4271: An Infrastructure to Support Secure Internet Routing. Technical report, 2012.

[97] M. Lepinski and K. Sriram. RFC 8205: BGPsec Protocol Specification, 2017.

[98] V. Li, F.and Paxson. A Large-Scale Empirical Study of Security Patches. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pages 2201–2215, 2017.

[99] X. Li, B. Liu, X. Zheng, H. Duan, Q. Li, and Y. Huang. Fast IPv6 Network Periphery Discovery and Security Implications. In 2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), pages 88–100. IEEE, 2021.

[100] Linksys. Differences between IPv4 and IPv6 (2020). Available from: https://www.linksys.com/us/support-article/?articleNum=139604. Last Accessed: 18 Jun 2020.

[101] Linksys. Configuring the MAC Filter Feature of the Linksys Smart Wi-Fi Router using the Local Access Interface. Available from: https://www.linksys.com/us/support-article?articleNum=143602., 2019. Last Accessed: 2 February 2020.

[102] R. Lippmann, D. Fried, et al. Evaluating Intrusion Detection Systems: The 1998 DARPA Off-line Intrusion Detection Evaluation. In Proceedings DARPA Information Survivability Conference and Exposition, volume 2, pages 12–26. IEEE, 2000.

[103] F. Loi and others. Systematically Evaluating Security and Privacy for Consumer IoT Devices. In Proceedings of the 2017 Workshop on Internet of Things Security and Privacy, pages 1–6, 2017.

[104] C. Lu. Overview of Security and Privacy Issues in the Internet of Things. Internet of Things (IoT): A vision, Architectural Elements, and Future Directions, pages 1–11, 2014.

[105] R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP Misconfiguration. ACM SIGCOMM Computer Communication Review, 32(4):3–16, 2002.

[106] R. Mahy, P. Matthews, and J. Rosenberg. RFC 5766: Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN). RFC 5766, April 2010.

[107] G. Maier, F. Schneider, and A. Feldmann. NAT Usage in Residential Broadband Networks. In International Conference on Passive and Active Network Measurement, pages 32–41. Springer, 2011.

[108] ManageEngine. ManageEngine OpUtils. Available from: https://www.manageengine.com/products/oputils/, 2024. Last accessed: 2024-01-18.

[109] T. Mattessich. Exploits and vulnerabilities of ip camera's. http://cysecure.org, 2012. Accessed: 2021-02-17.

[110] S. McConnell. Code Complete. Pearson Education, 2004.

[111] Microsoft. Support: IPv6 on Xbox One (2020). Available from: https://support.xbox.com/help/Hardware-Network/connect-network/ipv6-on-xbox-one. Last Accessed: 18 Jun 2020.

[112] J. Mirkovic and P. Reiher. A Taxonomy of DDoS Attack and DDoS Defense Mechanisms. ACM SIGCOMM Computer Communication Review, 34(2):39–53, 2004.

[113] MITRE. CVE-2016-7406. Available from: https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2016-7406, September 2016. Last accessed: 20 Oct 2020.

[114] MITRE. Common Vulnerabilities and Exposures. Available from: https://cve.mitre.org/., 2021. Last Accessed: 17 February 2021.

[115] A. Mitseva, A. Panchenko, and T. Engel. The state of affairs in BGP security: A survey of attacks and defenses. Computer Communications, 124:45–60, 2018.

[116] A. Mitseva, A. Panchenko, and T. Engel. The State of Affairs in BGP Security: A Survey of Attacks and Defenses. Computer Communications, 124:45–60, 2018.

[117] MongoDB. Mongodb. Available from: https://www.mongodb.com/, 2024. Last accessed: 2024-01-18.

[118] H. Moore. Security Flaws in Universal Plug and Play: Unplug. Don't Play. Rapid7 https://information.rapid7.com/rs/411-NAK-970/images/SecurityFlawsUPnP.pdf, 2013. Accessed: 2021-02-17.

[119] P. Morgner et al. Security Update Labels: Establishing Economic Incentives for Security Patching of IoT Consumer Products. arXiv:1906.11094, 2019.

[120] A. Muller, G. Carle, and A. Klenk. Behavior and Classification of NAT Devices and Implications for NAT Traversal. IEEE network, 22(5):14–19, 2008.

[121] A. Muller, N. Evans, C. Grothoff, and S. Kamkar. Autonomous NAT Traversal. In 2010 IEEE Tenth International Conference on Peer-to-Peer Computing (P2P), pages 1–4. IEEE, 2010.

[122] A. Muller, A. Klenk, and G. Carle. On the Applicability of Knowledge Based NAT-Traversal for Home Networks. In International Conference on Research in Networking, pages 264–275. Springer, 2008.

[123] C. Nader and E. Bou-Harb. Revisiting IoT Fingerprinting Behind a NAT. In 2021 IEEE Intl Conf on Parallel & Distributed Processing, pages 1745–1752. IEEE, 2021.

[124] M. Namestnikova. Router Security in 2021. Available from: https://securelist.com/router-security-2021/106711/, 2021. Last accessed: 13 August 2020.

[125] National Institute of Standards and Technology. Recommended Criteria For Cybersecurity Labelling For Consumer Internet of Things (IoT) Products. 2022. Last accessed 23 Oct 2022.

[126] Netgear. What is Explicit Beamforming and How Does It Work? Available from: https://kb.netgear.com/31299/What-is-explicit-beamforming-and-how-does-it-work, 2019.

[127] U. Networks. FAQ: Does AmpliFi Have a Firewall? (2020). Available from: https://help.amplifi.com/hc/en-us/articles/115009611867-Does-AmpliFi-have-a-firewall-. Last Accessed: 18 Oct 2020.

[128] M. Niemietz and J. Schwenk. Owning Your Home Nnetwork: Router Security Revisited. arXiv preprint arXiv:1506.04112, 2015.

[129] NIST. National Vulnerability Database. Available from: https://nvd.nist.gov/., 2021. Last Accessed: 17 February 2021.

[130] NIST. NIST RPKI Monitor. Available from: https://rpki-monitor.antd.nist.gov/ROV, 2023. Last accessed: 1 September 2023.

[131] N. Nthala and I. Flechais. Rethinking Home Network Security. 2018.

[132] K. Olson, J. Wampler, and E. Keller. Doomed to Repeat with IPv6? Characterization of NAT-centric Security in SOHO Routers. ACM Computing Surveys, 2023.

[133] K. Olson, J. Wampler, F. Shen, and N. Scaife. NATting Else Matters: Evaluating IPv6 Access Control in Residential Networks. In Proceedings of the Probability and Meaning Conference, 2021.

[134] Open Connectivity Foundation. UPnP+ Specification (2020). Available from: `https://openconnectivity.org/developer/specifications/upnp-resources/upnp/#upnp-plus`. Last accessed: 18 Oct 2020.

[135] J. Paillisse and others. IPchain: Securing IP Prefix Allocation and Delegation with Blockchain. In 2018 IEEE International Conference on Internet of Things (iThings), pages 1236–1243. IEEE, 2018.

[136] A. Pankratov. Server-Mediated Setup and Maintenance of Peer-to-Peer Communications, Oct. 23 2012. US Patent 8,296,437.

[137] M. Paul. IoT Security Flaw Leaves 496 Million Devices Vulnerable At Businesses: Report, 2019. Available from: `https://www.crn.com/news/internet-of-things/300106806/iot-security-flaw-leaves-496-million-devices-vulnerable-at-businesses/-report.htm`.

[138] R. Penno et al. RFC 9099: Operational Security Considerations for IPv6 Networks. RFC Draft Ver 21, April 2019.

[139] R. Penno and others. RFC 7857: Updates to Network Address Translation (NAT) Behavioral Requirements. BCP 127, April 2016.

[140] M. Petit-Huguenin, G. Salgueiro, J. Rosenberg, D. Wing, R. Mahy, and P. Matthews. RFC 8489:Session Traversal Utilities for NAT (STUN). RFC 8489, February 2020.

[141] J. Postel. RFC 792: Internet Control Message Protocol. STD 5, September 1981.

[142] M. Powell. Wi-Fi Router Security Knowledge Gap Putting Devices and Private Data at Risk in UK Homes. Available from: `https://www.broadbandgenie.co.uk/blog/20180409-wifi-router-security-survey`, 2018. Last accessed: 5 November 2021.

[143] A. Press. No Rush to Upgrade Your WiFi Router. Available from: `https://www.law.com/legaltechnews/almID/1167214009597/?id=1167214009597?id=1167214009597&slreturn=20211017152532`, 2006. Last accessed: 5 November 2021.

[144] P. Richter, F. Wohlfart, N. Vallina-Rodriguez, M. Allman, R. Bush, A. Feldmann, et al. A Multi-perspective Analysis of Carrier-grade NAT Deployment. In Proceedings of the 2016 Internet Measurement Conference, pages 215–229, 2016.

[145] R. I. E. (RIPE). Routing Information Service. Available from: `https://www.ripe.net/analyse/internet-measurements/`, 2024. Last accessed: 2024-01-18.

[146] J. Rosenberg. Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols. RFC 5245, April 2010.

[147] Rosenberg, J. and Weinberger, J. and Huitema, C. and Mahy, R. RFC 3489: STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs). RFC 3489, March 2003.

[148] T. Rowan. Negotiating WiFi Security. Network Security, 2010(2):8–12, 2010.

[149] F. Ruiz-Aliseda and P. Zemsky. Adoption is not Development: First Mover Advantages in the Diffusion of New Technology. INSEAD Business School, (2007/03), 2006.

[150] M. Saad, A. Anwar, A. Ahmad, H. Alasmary, and others. RouteChain: Towards Blockchain-based Secure and Efficient BGP Routing. Computer Networks, 217:109362, 2022.

[151] J. Saltzer, D. Reed, and D. Clark. End-to-end Arguments in System Design. ACM Transactions on Computer Systems (TOCS), 2(4):277–288, 1984.

[152] O. Santos. The Evolution of Scoring Security Vulnerabilities: The Sequel. Available from: `https://blogs.cisco.com/security/cvssv3-study`, 2016. Accessed: 2021-02-17.

[153] B. Schlinker, I. Cunha, Y. Chiu, S. Sundaresan, and E. Katz-Bassett. Internet Performance from Facebook's Edge. In Proceedings of the Internet Measurement Conference, pages 179–194, 2019.

[154] S. E. (SEED). SEED Internet Emulator. Available from: `Availablefrom:https://seedsecuritylabs.org/`, 2024. Last accessed: 2024-01-18.

[155] R. Shah and C. Sandvig. Software Defaults as De Facto Regulation the Case of the Wireless Internet. Information, Community & Society, 11(1):25–46, 2008.

[156] Y. Shavitt and E. Shir. DIMES: Let the Internet Measure Itself. ACM SIGCOMM Computer Communication Review, 35(5):71–74, 2005.

[157] Y. Shen and P. Vervier. Iot Security and Privacy Labels. In Privacy Technologies and Policy, pages 136–147. Springer, 2019.

[158] Y. Sheng, K. Tan, G. Chen, D. Kotz, and A. Campbell. Detecting 802.11 MAC layer spoofing using received signal strength. In IEEE INFOCOM 2008 - The 27th Conference on Computer Communications, pages 1768–1776, April 2008.

[159] J. Sherman. BGP Incident Data 2020. Available from: `https://www.atlanticcouncil.org/resources/bgp-incident-data/`, 2020.

[160] Shodan. UPnP Exposure Scan. Available from: `https://www.shodan.io/`. Last accessed 20 February 2021, 2021.

[161] H. Singh, W. Beebee, C. Donley, and B. Stark. RFC 7084: Basic Requirements for IPv6 Customer Edge Routers. RFC 7084, Internet Engineering Task Force, November 2013. Last Accessed: 18 Oct 2020.

[162] S. Singh, P. Sharma, S. Moon, and J. Park. Advanced Lightweight Encryption Algorithms for IoT Devices: Survey, Challenges and Solutions. Journal of Ambient Intelligence and Humanized Computing, pages 1–18, 2017.

[163] N. Skoberne, O. Maennel, I. Phillips, R. Bush, J. Zorz, and M. Ciglaric. IPv4 Address Sharing Mechanism Classification and Tradeoff Analysis. IEEE Transactions on Networking, 22(2):391–404, 2013.

[164] SMC Networks. SMC8014WG-SI User Manual. Available from: `https://manualmachine.com/smcnetworks/ezconnectsmc8014wgsi/479465-user-manual/`, 2015.

[165] M. Smith and M. Hunt. Network Security Using NAT and NAPT. pages 355–360, 2002.

[166] M. Smith and R. Hunt. Network Security using NAT and NAPT. In Proceedings 10th IEEE International Conference on Networks (ICON 2002), pages 355–360. IEEE, 2002.

[167] J. Snijders, W. Hargrave, K. Rechthien, M. Stucchi, and P. Hoogsteder. IXP Cost Comparison. Available from: http://peering.exposed, 2024. Last accessed: 2024-01-18.

[168] Y. Song, A. Venkataramani, and L. Gao. Identifying and Addressing Reachability and Policy Attacks in "Secure" BGP. ACM Transactions on Networking, 24(5):2969–2982, 2016.

[169] Srisuresh, P. and Holdrege, M. RFC 2663: IP Network Address Translator (NAT) Terminology and Considerations. RFC 2663, August 1999.

[170] SSC. Payment Card Industry Security Standards Council. Available from: https://www.pcisecuritystandards.org/, 2023. Last accessed: 13 August 2020.

[171] R. N. Staff. Ripe Atlas: A Global Internet Measurement Network. Internet Protocol Journal, 18(3):2–26, 2015.

[172] G. Starnberger. NAT Traversal Techniques in VoIP Protocols. PhD thesis, 2007.

[173] A. Steele. Keep Your Friends Close and Your Passwords Closer, 2016. Available from: https://blog.lastpass.com/2016/02/infographic-keep-your-friends-close-your-passwords-closer-2.html/.

[174] Synopsis. MiniUPnP. Available from: https://www.openhub.net/p/miniupnp, 2021. Last Accessed: 2021-02-17.

[175] P. Szewczyk and C. Valli. Insecurity by Obscurity: A Review of SoHo Router Literature from a Network Security Perspective. Journal of Digital Forensics, Security and Law, 4(3):1, 2009.

[176] K. Tabassum, A. Ibrahim, and S. El Rahman. Security Issues and Challenges in IoT. In 2019 International Conference on Computer and Information Sciences (ICCIS), pages 1–5. IEEE, 2019.

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

[178] The American Consumer Institute. Securing IoT Devices: How Safe Is Your Wi-Fi Router? Available from: https://www.theamericanconsumer.org/wp-content/uploads/2018/09/FINAL-Wi-Fi-Router-Vulnerabilities.pdf, 2018. Last accessed: 18 Oct 2020.

[179] ThousandEyes. ThousandEyes - Digital Experience Monitoring. Available from: https://www.thousandeyes.com/, 2024. Last accessed: 2024-01-18.

[180] Trendmicro. UPnP-enabled Home Devices and Vulnerabilities. Available from: https://www.trendmicro.com/en_us/research/19/c/upnp-enabled-connected-devices-in-home-upnpatched-known-vulnerabilities.html, 2019. Last Accessed: 2021-2-1.

[181] Tripwire. SOHO Wireless Router (In)Security. Available from: http://www.properaccess.com/docs/Tripwire_SOHO_Router_Insecurity_white_paper.pdf, 2014. Last accessed: 20 Oct 2020.

[182] K. Tykhonov. Blueborne Vulnerabilities in Bluetooth Implementations in Different Operation Systems, 2018.

[183] L. Ubiedo, T. O'Hara, M. Erquiaga, and S. Garcia. Current State of IPv6 Security in IoT. arXiv preprint arXiv:2105.02710, 2021.

[184] United States Code. Sarbanes-Oxley Act of 2002, PL 107-204, 116 Stat 745. Available from: https://www.bibsonomy.org/bibtex/21f18ad23845c0489c3002d950d8c2a27/dret, July 2002.

[185] G. Van de Velde, T. Hain, R. Droms, B. Carpenter, and E. Klein. RFC 4864: Local Network Protection for IPv6. RFC 4864, Internet Engineering Task Force, May 2007. Last Accessed: 18 Oct 2020.

[186] M. Vink, E. Poll, and A. Verbiest. A Comprehensive Taxonomy of Wi-Fi Attacks. PhD thesis, Radboud University Nijmegen Nijmegen, The Netherlands, 2020.

[187] T. Wan, E. Kranakis, and P. van Oorschot. Pretty Secure BGP, psBGP. In NDSS, 2005.

[188] D. Weber. A Taxonomy of Computer Intrusions. PhD thesis, Massachusetts Institute of Technology, 1998.

[189] P. Weidenbach and J. vom Dorp. Home Router Security 2020. Available From: https://www.fkie.fraunhofer.de/content/dam/fkie/de/documents/HomeRouter/HomeRouterSecurity_2020_Bericht.pdf, 2020. Last Accessed: 2021-02-17.

[190] R. White. Securing BGP through Secure Origin BGP (soBGP). Business Communications Review, 33(5):47–53, 2003.

[191] C. Wimmer. Wireless LAN Security in a SOHO Environment: A Holistic Approach. GRIN Verlag, 2012, 2008.

[192] D. Wing and others. RFC 6887: Port Control Protocol (PCP). RFC 6887, Internet Engineering Task Force, April 2013. Last Accessed: 18 Oct 2020.

[193] J. Woodyatt. Rfc 6092: Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service. RFC 6092, Internet Engineering Task Force, January 2011.

[194] A. Wool. A Quantitative Study of Firewall Configuration Errors. Computer, 37(6):62–67, 2004.

[195] Q. Xing, B. Wang, and X. Wang. Bgpcoin: Blockchain-based Internet Number Resource Authority and BGP Security Solution. Symmetry, 10(9):408, 2018.

[196] J. Xiong and K. Jamieson. Securearray: Improving WiFi Security with Fine-Grained Physical-Layer Information. In Proceedings of the 19th Annual International Conference on Mobile Computing & Networking, pages 441–452, 2013.

[197] L. Zhang. A Retrospective View of Network Address Translation. IEEE Networks, 22(5):8–12, 2008.

[198] L. Zhang, W. Jia, et al. Research of TCP NAT Traversal Solution Based on Port Correlation Analysis & Prediction Algorithm. In 2010 6th International Conference on Wireless Communications (WiCOM), pages 1–4. IEEE, 2010.

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