Doomed to Repeat with IPv6? Characterization of NAT-centric Security in SOHO Routers

KARL OLSON, JACK WAMPLER, and ERIC KELLER, University of Colorado, Boulder, Colorado, USA

With the transition to IPv6, addressing constraints that necessitated a common security architecture under network address translation (NAT) are no longer present. Instead, manufacturers are now able to choose between an open model design, where devices are end-to-end reachable, or a more familiar closed model, where the home gateway may continue to serve as a perimeter security device. The potential for further nuance, such as differences in default access control policies, filtering behaviors, and IPv6 specific requirements, present an environment defined by ambiguity. For the consumer, the potential impact of these changes are unclear. To address this uncertainty, we taxonomize the present NAT-centric model of consumer gateway security through a survey of over 300 common vulnerabilities and exposures surrounding NAT and hole punching protocols. From this survey, we contextualize the limited security NAT has provided while serving as the primary perimeter defense mechanism in home networks. We further define how this baseline security model for consumer gateways is reflected in IPv6 through an assessment of ten commonly deployed consumer gateways. Our conclusion is that familiarity of a NAT-centric design is no longer assured for IPv6, requiring an active involvement by users to limit exposures within their home networks.

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Networks** → **Network design principles**; • **Security and privacy** → **Network security;**

Additional Key Words and Phrases: Network address translation (NAT), hole-punching, consumer gateway security, IPv6, device standards

ACM Reference format:
Karl Olson, Jack Wampler, and Eric Keller. 2023. Doomed to Repeat with IPv6? Characterization of NAT-centric Security in SOHO Routers. *ACM Comput. Surv.* 55, 14s, Article 305 (July 2023), 37 pages.
https://doi.org/10.1145/3586007

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)** toward 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 [14]. In the context of residential networks, that value jumps to

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

Authors' address: K. Olson, J. Wampler, and E. Keller, University of Colorado, 1111 Engineering Dr. Boulder, CO 80309; emails: {kaol6371, jack.wampler, eric.keller}@colorado.edu.

0360-0300/2023/07-ART305 $15.00
https://doi.org/10.1145/3586007

nearly 95% [70]. Without the widespread deployment of NAT, the Internet could not support the 40 billion connected devices today [39].

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 [31, 45, 108, 128]. 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 [30, 93, 117] 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 [100].

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 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 to better understand and define requirements for IPv6. We conclude by answering how manufacturers are, at present, 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 that 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 article is organized as follows: In Section 2, we define the common security properties, operational models, supporting parties, and attacker goals to provide a common understanding of the complex interrelationships involved in defining a common access control

model for consumer gateways. In Section 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 4 and 5. We use this context to taxonomize the operational failures of both NAT and hole-punching methods in Section 6 and conclude with a trend analysis to show that consumer gateway security has never been great in Section 7. We use these efforts to identify pitfalls and requirements for securing IPv6 consumer gateways in Section 8. Finally, we provide key takeaways and recommendations for improving the default access control model for consumer gateways operating IPv6 in Section 9. A review of related work concludes our survey in Section 10.

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. 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.

2.1 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, as shown in Figure 1 (Left).

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, as shown in Figure 1 (Right). This open model approach is commonly found with 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 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,

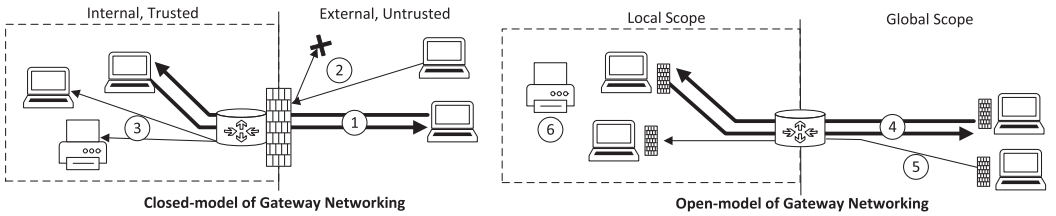

Fig. 1. Two architectural approaches to gateway networking. In the closed model (left image), a gateway acts as the primary device to provide access control into a local network. (1) Network communications from the internal network are allowed outbound with connection state maintained to match and allow 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.

protocol, or architecture should adhere. We briefly define these core security properties 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 [77].

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 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 that could further system exposure.

Reliability, Authenticity, and Non-Repudiation are recommended extensions of the CIA triad by ISO 2700 that further define how security goals may be achieved [34]. 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 that 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.

2.3 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 demonstrate and discuss further in Sections 8 and 9.

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. However, 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.

2.4 Attacker Goals

Finally, 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 References [61, 68, 121]. 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 that 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.

2.5 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, 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

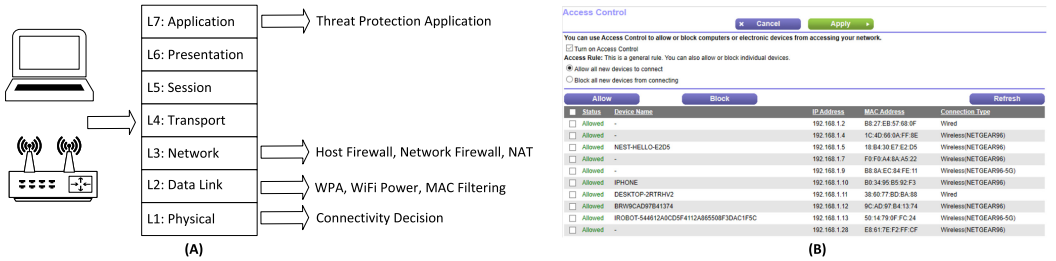

Fig. 2. Home environment access control methods. (A) An OSI layer view of typical access control methods available to consumers. (B) Home systems and device manufacturers do little to ease identification of systems. While some devices show enough information to identify, many do not, making it hard to assess devices on a network for most users.

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.

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) [17, 18]. 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 [36, 124]. 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 policy, consumers are left to presume that their connected devices do not behave maliciously, though evidence shows this assumption to consistently fail in practice [71, 72, 118].

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 and discuss in detail below.

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 that complicate low-level access control. The nature of **Radio Frequency (RF)** transmissions in the 2.4 and 5 GHz bands means that they frequently leak beyond the bounds of the physical location of the transmitter [98]. An adversary in an off-site location can collect these signals, disrupt, or attempt to connect to the network. While beam-forming [83] and secure arrays [127] can alleviate these issues, the user must still monitor the network for unauthorized devices. If unauthorized devices are found on the wireless network, then

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.

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 [67, 84]. 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 [44, 103], 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(b) 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 [5]. 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 [3] 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 [112]; some platforms (e.g., iOS and Windows) provide features that allow user to automatically share a wireless password with a nearby contact [9]. 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.

3.3 Layers 3 and 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 [126] 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 [93].* The popular traditional NAT-PT mode of NAT (described in Section 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 third party, often unbeknownst to the user [1, 37]. Restricting this permissive outbound behavior is challenging at best for reasons previously mentioned.

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. Last, 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 [115].

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 that show that users rarely involve themselves with changing default configuration settings or do so in a way that improves their security [59, 102].

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. 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 to glean lessons for IPv6 deployment.

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 [128]. 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.

4.0.1 Traditional NAT and NAPT. **Traditional NAT (NAT)** maintains a single external IP address that 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 3. At larger

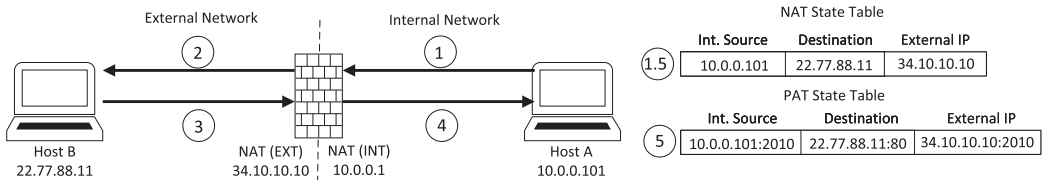

Fig. 3. Traditional NAT. (1) Host A initiates connection to host B reaching a NAT gateway. (2) NAT gateway maps Host A IP address to external globally routable address, updates the IP packet to reflect the external interface IP and forwards packet to Host B. A connection state table within the gateway 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.

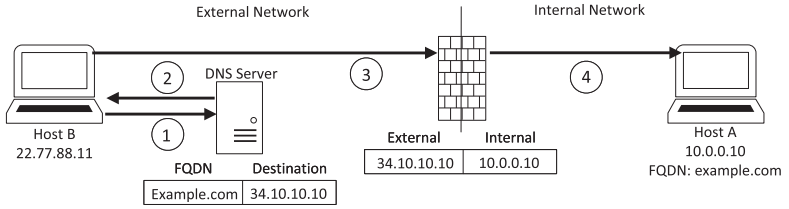

Fig. 4. Two-way NAT. (1) Host B seeks to establish communication with Host A, located behind a NAT gateway, by first querying the public DNS server for the FQDN and external IP address of a hosted service. (2) DNS server responds with the public IP associated with FQDN. (3) Host B sends request to public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

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 [110].

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 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 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 [110].

4.0.3 Twice-NAT. With Twice-NAT, both the source and destination address of a packet are translated, as shown in Figure 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 [110]. The concern is that an internal

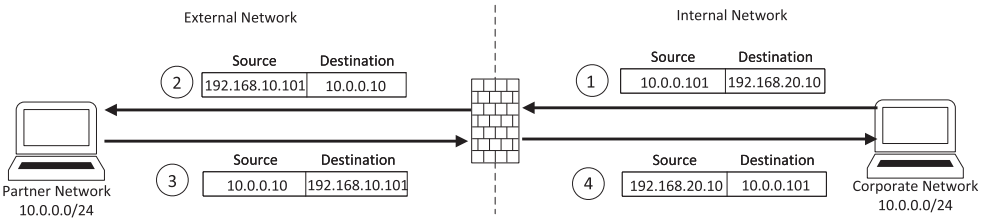

Fig. 5. Twice NAT. (1) Host on corporate network seeks to communicate with partner office where both locations operate the same private IP space. Internal to the corporate network, the external partner network is assigned as an alternate IP space to prevent internal routing conflicts. (2) When communication 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.

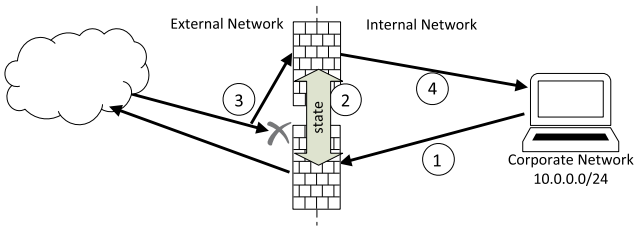

Fig. 6. Multi-homed NAT. (1) Host on corporate network utilizes NAT to reach an external system. (2) NAT gateway updates state table and synchronizes state across all gateways. (3) In the event of an outage involving the primary NAT gateway, return traffic defaults to secondary gateway. (4) The secondary gateway finds the synchronized mapping in state table and forwards traffic to appropriate host on internal network.

host may have the same routable address as an external host. When communication is executed internally, the request will not make it 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.

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 6 demonstrates a typical configuration in multi-homed NAT networks. Here, gateway #1 may be the primary NAT path that shares state information with gateway #2. In case of a failure all traffic is rerouted to gateway #2 transparently, ensuring communication is uninterrupted.

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 on forwarding and response characteristics employed by gateway manufacturers [97], see Figure 7:

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

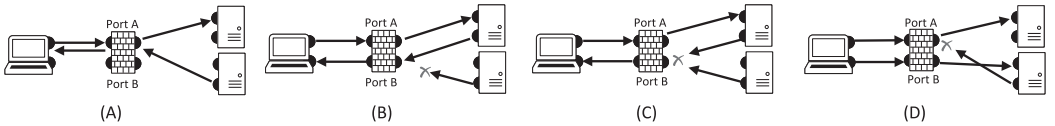

Fig. 7. NAT operational methods. (A) Full-cone NAT. (B) Address-restricted Cone NAT. (C) Port-restricted Cone NAT. (D) Symmetric NAT.

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 [81].

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 [81].

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 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 [81].

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, then 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 [81], despite presenting the strongest assurance for access control.

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 [2, 43, 57]. 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 [43].

Often these response characteristics are undocumented, requiring a consumer to conduct detailed testing of their gateway 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 that address the many use cases [79]. 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 7, this is rarely achieved in practice.

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 that 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) [10, 41, 42, 110]. This process of continual refinement continues with the most recent publication of NAT behavior requirements published in 2016 under RFC 7857 [32].

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 [10]. If a user had sufficient technical understanding, then 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 [89]. 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.

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.

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 that allow for seamless and automated gateway configuration, as shown in Figure 8. Here, a gateway daemon listens for local network participants to execute a configuration

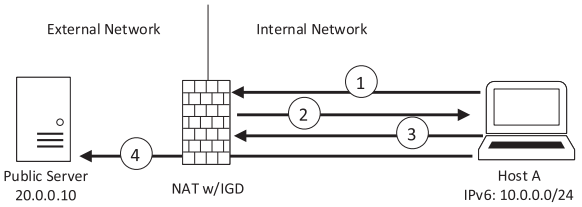

Fig. 8. UPnP with IGD. (1) A new device executes a Simple Service Discovery Protocol (SSDP) request to identify supporting devices on the local network. Identified devices respond with a location (e.g., 192.168.1.1/service.xml) for the host to find defined services available. (2) Host request services 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 [11].

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 [4, 13, 77]. 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 [124]. 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.

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 that 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 [35]. With UDP hole-punching, a publicly accessible server acts as a mediator to coordinate connection establishment, as shown in Figure 9. If the connection is dropped, then 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 [46].

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 [92]. Due to the autonomous nature, setup requires strict coordination, prior knowledge of the endpoint gateway IP address, and predictable port selection to successfully execute [52, 92]. While shared knowledge of the destination IP address is easy to coordinate, the shared knowledge of which port will be generated is not [52]. Depending on vendor implementation of NAT, the selection of a port may be predictable using a

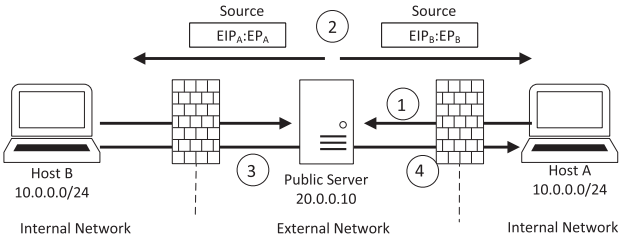

Fig. 9. UDP hole-punching. (1) Both host A and B establish communication with well-known public proxy about their intent to connect, resulting the each gateway establishing a port mapping back to the internal host, e.g., $extPort_A : intIP_A$. (2) The public proxy server inspects both communication streams 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.

simple known algorithm, direct mapping, or sequential selection [129]. If these port determination methods are not predictable, then 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 [35]. 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 [80].

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 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% [90]. In this server mediated method, each host initially establishes communication with an external moderator, as shown in Figure 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

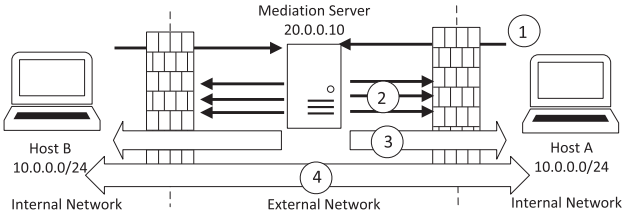

Fig. 10. Hamachi tunneling. (1) Both host A and B establish communication with well-known mediation server through local NAT gateways. (2) The mediation server directs each host to execute a series of probes to learn about each gateway's forwarding characteristic and port assignment process. (3) The mediation server begins tunnel setup to each host and monitors for success. (4) If tunnel setup successful, then the mediation server provides each endpoint with the other party's $extIP:extPort$ information and hands off tunnel so each host may communicate directly.

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 [53].

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)** [69, 91, 96], **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 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 [60].

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.

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 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

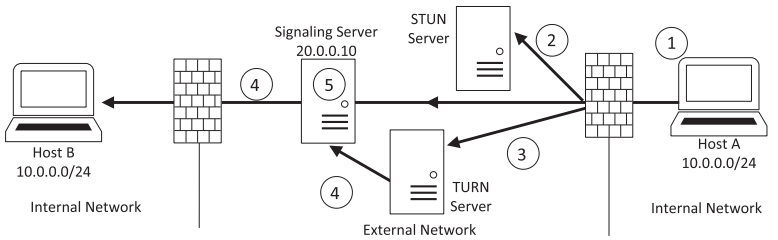

Fig. 11. ICE with TURN and STUN. (1) Each host will go through a discovery process of learning all potential usable addresses in which to communicate with a remote peer. These candidate addresses will typically include connected interfaces (physical, virtual, or tunnel), (2) public facing gateway addresses discovered through STUN, or (3) if a TURN server is specified for relaying communications, then 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.

> kind of firewall, any NAT gateway implementing this protocol SHOULD have an administrative mechanism to disable it, thereby restoring the pre-NAT-PMP behavior [16].

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." Second, 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 [55, 93]. 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.

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 [33]. Recent efforts to assess IPv6 operation within consumer gateways present many of the same challenges and pitfalls [89]. The result is an ambiguous operational environment of technologies within consumer gateways that have no clear security or operational guarantee.

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. To align our efforts with previous works, we adopt the literature review and organizational methods utilized in References [63, 75].

For our review, we rely on the NIST NVD [86], the MITRE CVE [76], and the **U.S. Computer Emergency Readiness Team (US-CERT)** Vulnerability Notes databases. We performed a vulnerability search and selection process over five steps, demonstrated in Figure 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

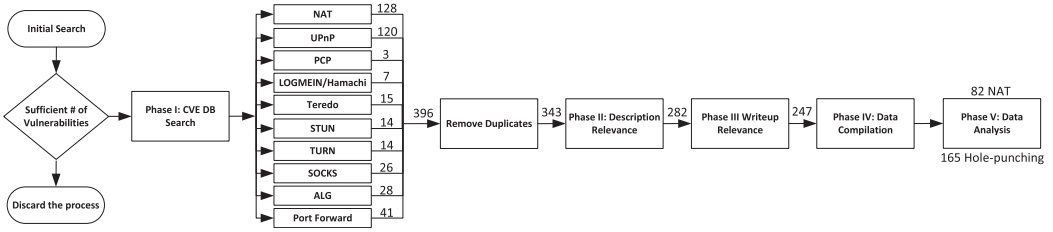

Fig. 12. CVE review process map. Relevant hole-punching protocol technologies were searched and sorted to assess overall takeaways. In total, 82 CVEs related to NAT and 165 related to hole-punching methods were used in our assessment.

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.

6.1 Hole-Punching and NAT Security Taxonomy

The CVE 2.0 categorizations provided by the 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 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 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.

Below, we present our taxonomy classification categories, along with related statistics resulting from our analysis, which we use to draw security takeaways at the conclusion of this section.

6.2 Taxonomy Category Classifications

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

[2]We expand our search beyond just consumer gateway security failures 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 that *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.

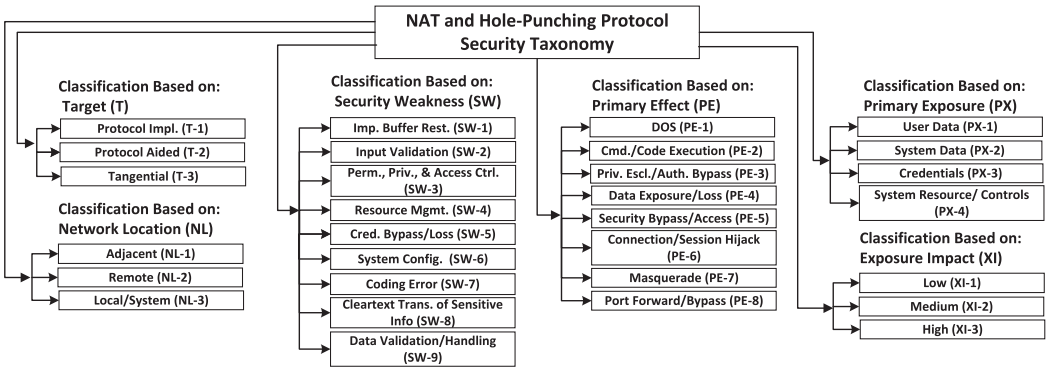

Fig. 13. Final taxonomy of hole-punching and NAT security failures. The taxonomy is based on systematic review of 300 CVE documents to obtain source classifications of security flaws based on target, network vector, security weakness, attack effect, and resulting exposure. A single taxonomy is presented to represent an overall assessment of gateway security flaws and due to the significant overlap found in conducting the taxonomies separately.

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 that 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 [20]. 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.

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 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 7.3.

Local/System (NL-3) attacks require direct access to the target device 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.

6.2.3 Classification Based on Security Weakness (SW). 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 (e.g., 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 that 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 [21]. 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 that 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 [77]. 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 IGD [27].

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 [22]. 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.

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. 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 [77]. 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 that 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 [24].

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 [25]. 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 [26]. 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 [23].

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.

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 that allowed a remote attacker to eavesdrop via publicly exposed STUN ports [73].

System Data (PX-2) exposure consists of device information such as type, configuration, or protocol communication traffic that 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.

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.

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 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 [47]. 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 [99].

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.

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 [94, 99]. 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 14 shows that over the life-cycle of hole-punching methods there has been a steady rise in discovered vulnerabilities. This

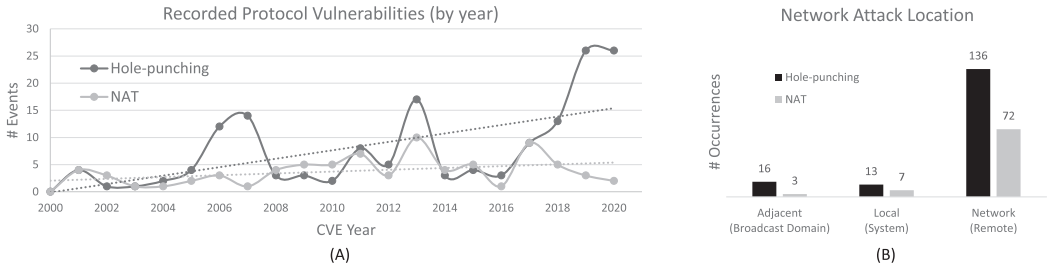

Fig. 14. CVE distribution by year and primary network avenue of attack. (A) Documented vulnerabilities for hole-punching methods have increased over time, in line with broader security trends overall. Contrary to this increase, NAT has remained relatively unchanged, averaging 1.8 CVE's per year. Assessment 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.

growth is in line with general trends in CVE reporting overall [65]. 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 [122]. 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 [77].
(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 [113]. With an average of fifteen to fifty bugs per one-thousand LoC, the potential for mistakes in hole-punching packages increases significantly [74].
(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 [116].
(4) Updating separate software packages can be costly from a development perspective, as changes may introduce second- and third-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.

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 [101]. 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

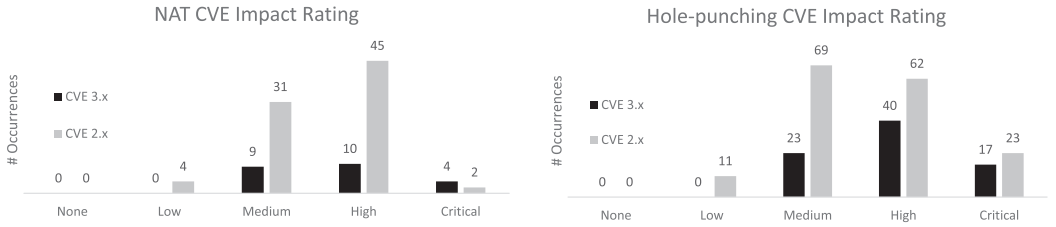

Fig. 15. CVE scoring distribution—NAT and hole-punching protocols. Scores for CVE 2.x and 3.x are displayed. Scoring for 2.x represents *all* vulnerabilities while the 3.x reflects only the CVEs that have been scored under the newer metric. The CVE 2.x scoring is aligned to the 3.x categories based on 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.

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 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 [117], 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 [7, 40]. 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 [47].

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 that 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 14(b). 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 [28]. Similar examples exist throughout the network vector classifications.

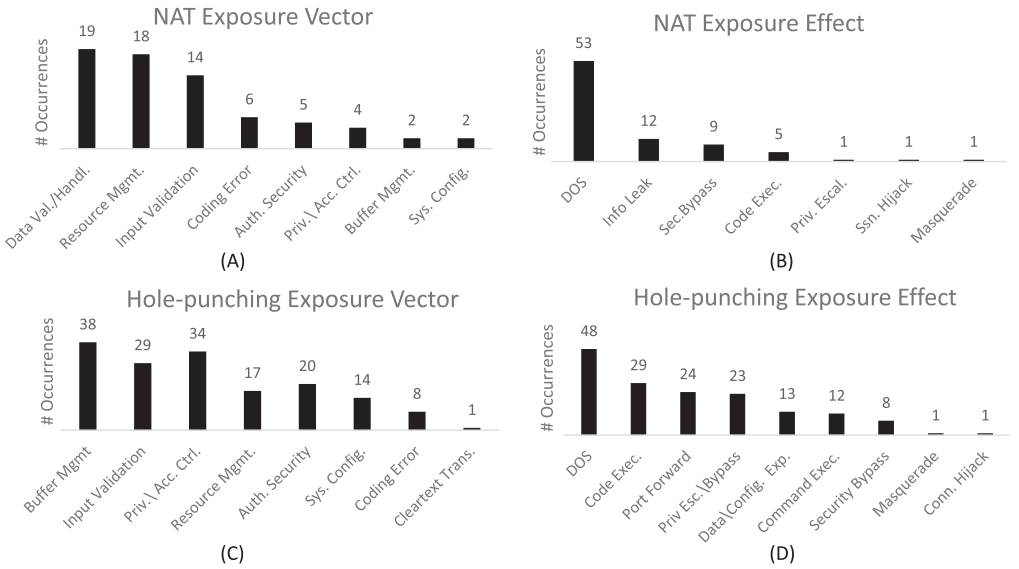

Fig. 16. Enumerated weakness vectors and effects. Enumerated occurrences of weakness vectors and resulting effects within NAT and hole-punching methods are presented. Note: hole-punching enumeration does not include six undisclosed vulnerabilities that could not be categorized.

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 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 that allowed access to internal networks [77]. Globally, nearly 450,000 devices still maintain this exposure eight years later [104].

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

7.5 CIA Triad Exposure

Within the CIA Triad, user exposures are significantly higher within hole-punching methods, as demonstrated in Figure 17(a); 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 [67, 77].

Contrary to these extreme exposures within hole-punching methods are the exposures related to NAT, shown in Figure 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

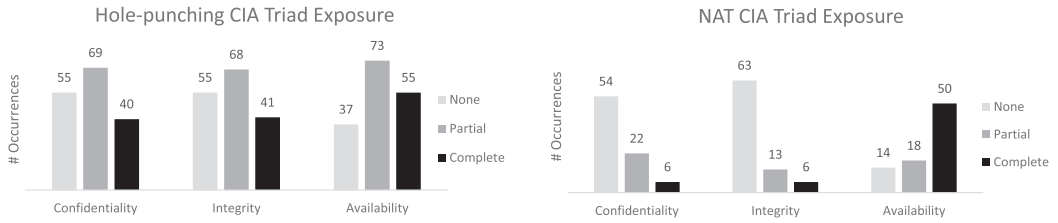

Fig. 17. CIA triad exposure—Hole-punching methods (left) and NAT (right). Exposures to confidentiality, integrity, and availability of systems due to exposures within or created by either NAT or hole-punching methods are presented. Exposures are ranked on a partial, complete, and none categorization representing the degree of exposure from an attack.

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.

8 POST NAT - DEFAULT IPV6 SECURITY POLICIES AND ACCESS CONTROL MECHANISMS

While IPv4 NAT provides connectivity for a majority of consumer devices today, many ISPs are actively integrating and transitioning to IPv6. What this means for consumer security is unclear. The potential for an open model architecture is a strong departure from the closed model operation of NAT, which consumers are familiar with. While not a security flaw by definition, this departure is a paradigm shift in gateway operation. Here, users must be made aware of and understand the need to implement additional security methods, such as endpoint security or carefully managed firewall policies. Further, in absence of NAT, it is unclear what a default firewall security policy may look like or how manufacturers will implement controls to precisely capture IPv6 nuances, such as multi-homing, multiple addressing scopes, or dynamic address generation. Drawing from our prior work investigating IPv6 security within the home [89], we provide a short summary of our results to highlight that these challenges are indeed real.

8.1 IPv6 Assessment Process

For our assessment, we considered two overarching questions: What does a default access control policy for IPv6 look like and how can a consumer enact security policy change? To answer these two questions, we conducted a review ten consumer gateways as follows: (1) Obtain a representative sample of home gateways to perform an assessment on. (2) Perform a review IPv6 security characteristics and default configurations. Here, characteristics are defined as any element that may play a role in IPv6 security policy. (3) Validate operation of controls by conducting a series of scans from multiple vantage points to validate control response. We provide the details of each process below.

8.1.1 Obtaining a Representative Sample of Consumer Gateways for Analysis. 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 [64]. Of 4.8K router vendors globally, we selected 12 routers that 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 as marketed. The remaining ten devices used in our assessment are shown in Table 1.

8.1.2 Determine IPv6 Security Characteristics and Default Policies. To understand how a consumer may effect IPv6 security policy, we conducted a review of each gateway to identify IPv6 features, mechanisms, and characteristics that may have a role in determining the overall security of a consumer network. Upon completing this review, we arrived at the following device characteristics and considerations for our assessment:

- **Default IPv6 Routing.** We first checked if each router supports IPv6 and whether it enables that support by default.
- **Firewall Present.** Next, we evaluated 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, then we evaluated 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 [105]. We evaluated whether or not the device included this functionality.
- **Security Warning.** When the One-Click Open option was used, we evaluated if there was any warning or communication to the user about the danger of disabling the firewall.
- **Rule Generation.** We evaluated whether each device included the ability to create exceptions to the default firewall policy.
- **IP Specification.** In cases where a user is allowed to implement policy, we assessed whether or not they are able to specify policy by an IP.
- **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 further evaluated whether rules can be created by specifying a device (e.g., by MAC address or another identifier) rather than a specific IP address and whether this policy applied to the entire range of addresses for that device.
- **IPv6 UPnP Support.** Finally, we evaluated 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.

8.1.3 Device Policy Assessment and Validation. Since routers do not explicitly advertise their firewall policies, we conducted a series of black-box scans to establish the default filtering model, firewall filtering policies, hosted router services, and whether or not policy changes were correctly actioned. A complete assessment of each gateway involved nine total scans from two sources, each conducted with the firewall on and off as shown in Figure 18. The combination of sources and targets allowed complete measurement of IPv6 filtering policies, exposure, and default operational model of the gateway. These results were then compared with our evaluation of basic router characteristics to arrive at our final assessments of how IPv6 security is being implemented in the absence of a default-deny policy necessitated by NAT.

8.2 Results

In general, we found customer gateways with IPv6 capability have little commonality for baseline behaviors and policy mechanisms across the various manufacturers. Table 1 presents an overview of our findings showing a wide variance in access control policy, device security, and user security policy mechanisms for IPv6.

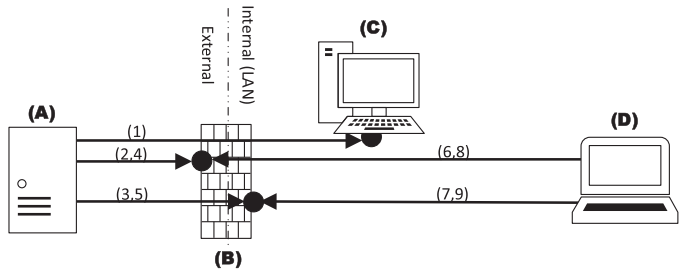

Fig. 18. Scanning protocol. To fully evaluate the security policy of each router, we scan from two vantage points (A) and (D) against three targets: (C) an internal host and (B) the firewall internal and external interfaces. In total, we conducted nine unique scans for each router as follows: Scan (1) 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; scan (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 (scans (4) and (5)). This process is repeated from the internal vantage first targeting the exposed services on each router interface (scans (6) and (7)) before conducting the same banner grab on exposed services (scans (8) and (9)).

Table 1. Routers—Heterogeneous Nature of Management Options and Default Configurations among the Devices Evaluated

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	●	○	–	–	○	○	–	–	○

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.

8.2.1 Open vs. Closed Model of Security. Two of the ten routers implemented the open-model architecture for security as the default IPv6 security policy (TP-Link AC1750 and Motorola MR2600). Both devices further enabled IPv6 routing under the default configuration. Neither device communicated this design decision to the consumer. In the case of the TP-Link, no implementation of a firewall for IPv6 was present despite allowing for IPv6 routing. In this case, not only was a customer exposed, there was little they could do about it absent actively involving themselves and turning off IPv6 routing altogether. When assessed, the TP-Link AC1750 was Amazon US's top-selling router [6] and TP-link was the top global provider, accounting for 15.9% of all deployed devices [64]. This suggests that a default open model may be commonly deployed in practice.

In accordance with RFC 7368, manufacturers are to provide a mechanism by which a customer may toggle between the open or closed security models for IPv6 [12]. In our assessment, only six of the ten gateways did so. Oddly, in only one instance did any device provide a warning to the user on the impacts of enabling the open model (e.g., disabling the firewall) for IPv6. 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 now creates.

8.2.2 Security Control Mechanisms. 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 altogether.

For routers that provide an interface to create exceptions to the default firewall filtering policy (pinholes), we found that two of six connect IPv6 policy rules to a device based on MAC address only. Unfortunately in both of these cases, traffic destined for *any* associated IPv6 address for the device is forwarded, providing no security value whatsoever. We suspect this was improper conversion of IPv4 filtering mechanisms, as the same controls were present for each device and filtered correctly for IPv4.

In the remaining four gateways users could provide a single, static address that the rule applies to. These rules were not updated if the device migrated or was assigned additional IPv6 addresses over time. This places a heavy burden on the user to be aware of their individual device addresses while also having to consider how these may change over time.

In one case, the Ubiquiti AmpliFi router, users were presented with an automatic pinholing solution through UPnP for IPv6. While it is unclear at present how many devices may utilize this configuration feature, in the absence of a device supporting UPnP for IPv6, a customer is presented with a binary option of disabling the firewall altogether as the only policy configuration mechanism.

8.2.3 Default Filtering Policies. We find that when CE routers are globally accessible a majority of them expose open services to the Internet as shown by Table 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.

9 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 to deliver 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 [117]. 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.

9.1 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

Table 2. Externally Exposed Services—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

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, **2,601**, 1,080, **10,000**
Linksys EA6350 AC1200	●	—	25, 53, **80**, 135, 139, **443**, 445, **2,601**, 1,080, **10,000**
Motorola MR2600	○	25, 135, 139, 445, 1,080	25, 135, 139, 445, 1,080
Nighthawk X4 R7000	●	—	25, 43, 80, 135, 139, 443, 445, 548, 1,080, **2,601**
Surfboard SBG10 DOCSIS 3.0	●	—	25, 80, 135, 139, 443, 445, 1,080
TP-Link AC1750 v2	○	No Enable Option	**22**, 25, 135, 139, 445, 1,080

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.

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 [125]. 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 [33]." 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 [29]. The end result is more in line with "assured exposure" instead of a secure default for the non-involved user.

9.2 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 [29, 66]. Implementation of a IPv6 stateful filter similar to NAT would help address many of these challenges.

9.3 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 [13, 38]. Despite ten years of efforts advising consumers to disable this feature, UPnP remains an on-going challenge [104]. 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 [47].

9.4 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 [48, 106]. As a result, many devices overwhelmingly rely on simple local authentication methods as the only means for access control [48]. 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 [119]. While these protocols do provide for organic security measures, such as encryption and authentication, many devices cannot support the computational overhead introduced [106]. 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 [48]. The result is both an increased need and expectation for consumer gateways to precisely provide default security assurances and mechanisms by which to adjust.

9.5 Incentives to Address Security

Historically, home router security has never been good [87]. Systems have been found to implement software packages over ten years old and in many cases never fix exposures at all [47]. 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. Solutions, such as device labelling

standards or contractual support guarantees, have been proposed as possible mechanisms to incentivize better security through improved consumer awareness [78, 88]. 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 as a necessary means has already occurred as demonstrated by California enacting state law SB-327 in 2017. This law required manufacturers to no longer use default credentials on devices sold within the state [56]. The law further requires that mechanisms must be present to force a user to create a unique password on initial device setup. However, more regulation is likely necessary to force broader security changes within the consumer gateway industry.

10 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 References [10, 32, 41, 42, 51, 110]. These are supported by academic efforts that further identify and define device specific behaviors related to unique vendor implementations [43, 49, 79] and the IETF's response to help standardize these behaviors [32]. Considerations for specific architectures include ISPs [95, 107] and home networks [49, 79]. 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) [15, 19]. 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 [35] and VoIP[54, 111] 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 that consider the role of NAT as a security mechanism are first discussed in Reference [109] with many subsequent works [8, 58, 114] 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 [128].

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 [62] and wireless [120, 123] both provide recent a detailed analysis of access control mechanisms and attacks surrounding WiFi. Authors of Reference [85] 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 [50, 82].

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.

11 CONCLUSION

After 30 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. The addition of mechanisms meant to further abstract consumer involvement in establishing refined security policy, however, demonstrates 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.

Looking forward, we see the underpinnings of many of the same mistakes occurring, particularly with IPv6 deployment. At present, open-ended requirements for IPv6 operation within consumer gateways enables a security environment defined by ambiguity. 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 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. Manufacturers, in particular, have a much larger responsibility, as many security flaws can be directly attributed to poor implementation or maintenance practices. 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 and that they are slow to patch, if ever.

Absent stronger standards or policies, or a more involved consumer, there is little here to look forward to with regards to consumer gateway security moving forward.

REFERENCES

[1] Abdifatah Abdi-Nur. 2017. Smart TV upgrade, privacy downgrade? *J. Colloq. Info. Syst. Secur. Edu.* 5, 1, 22–22.

[2] Bernard Aboba and William Dixon. 2004. *IPsec-Network Address Translation (NAT) Compatibility Requirements.* RFC 3715.

[3] Anne Adams and Martina Sasse. 1999. Users are not the enemy. *Commun. ACM* 42, 12 (1999), 40–46.

[4] Akamai. 2018. *UPnProxy: Blackhat Proxies via NAT Injections.* Technical Report.

[5] WiFi Alliance. 2022. WiFi Alliance Wireless Specifications. Retrieved from https://www.wi-fi.org/discover-wi-fi/specifications. Accessed: 2022-11-17.

[6] Amazon.com. 2020. Amazon Sales Popularity—Computer Routers (2020). Retrieved from https://web.archive.org/web/20201023233343/https://www.amazon.com/gp/bestsellers/pc/300189/ref=zg_b_bs_300189_1. Last accessed 23 October 2020.

[7] Amazon.com. 2021. Amazon EERO Technical Specification. Retrieved from https://support.eero.com/hc/en-us/articles/209962973-Frequently-asked-security-questions. Last Accessed: 1 February 2021.

[8] Cedric Aoun and Elwyn Davies. 2007. *Reasons to Move the Network Address Translator-Protocol Translator (NAT-PT) to Historic Status.* Technical Report. RFC 4966.

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

[10] Francois Audet and Cullen Jennings. 2007. *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP.* BCP 127.

[11] Mohamad Boucadair, Reinaldo Penno, and Dan Wing. 2013. *Universal Plug and Play (UPnP) Internet Gateway Device—Port Control Protocol Interworking Function (IGD-PCP IWF).* RFC 6970.

[12] A. Brandt, Sigma Designs, O. Troan, J. Weil, and Time Warner Cable. 2014. RFC 7368: IPv6 home networking architecture principles.

[13] Yunus Cadirci. 2020. *CallStranger Technical Report.* Technical Report. https://github.com/yunuscadirci/CallStranger.

[14] Martin Casado and Michael Freedman. 2006. Illuminating the shadows: Opportunistic network and web measurement. Retrieved from http://illuminati.coralcdn.org/stats.

[15] Whai-En Chen, Ya-Lin Huang, and Han-Chieh Chao. 2008. NAT traversing solutions for SIP applications. *EURASIP J. Wireless Commun. Netw.* 2008 (2008), 1–9.

[16] S. Cheshire and M. Krochmal. 2013. RFC 6886: Nat port mapping protocol (NAT-PMP). *IETF* (2013).

[17] Cisco. 2020. Cisco Meraki-go: Easy Networking for Busy People. Retrieved from https://www.meraki-go.com/.

[18] Frontier Communications. 2020. Frontier Home Internet Setup Guide. Retrieved from https://frontier.com/helpcenter/topics/install-fiber-optic Last Accessed: 10 March 2020.
[19] M. Aurel Constantinescu, V. Croitoru, and D. Oana Cernaianu. 2005. NAT/firewall traversal for SIP: Issues and solutions. In *Proceedings of the International Symposium on Signals, Circuits, and Systems (ISSCS'05)*, Vol. 2. IEEE, 521–524.
[20] MITRE Corporation. 2006. CVE-2006-2559. Retrieved from https://cve.mitre.org.
[21] MITRE Corporation. 2012. CVE-2012-0383. Retrieved from https://cve.mitre.org.
[22] MITRE Corporation. 2013. CVE-2013-3182. Retrieved from https://cve.mitre.org.
[23] MITRE Corporation. 2013. CVE-2013-6949. Retrieved from https://cve.mitre.org.
[24] MITRE Corporation. 2017. CVE-2017-17746. Retrieved from https://cve.mitre.org.
[25] MITRE Corporation. 2017. CVE-2017-7405. Retrieved from https://cve.mitre.org.
[26] MITRE Corporation. 2020. CVE-2020-16894. Retrieved from https://cve.mitre.org.
[27] MITRE Corporation. 2020. CVE-2020-25988. Retrieved from https://cve.mitre.org.
[28] MITRE Corporation. 2021. CVE-2007-2390. Retrieved from https://cve.mitre.org.
[29] Jakub Czyz, Matthew Luckie, Mark Allman, Michael Bailey et al. 2016. Don't forget to lock the back door! A characterization of IPv6 network security policy. In *Proceedings of the Conference on Network and Distributed Systems Security (NDSS'16)*.
[30] Nicholas De Leon. 2019. Many Wireless Routers Lack Basic Security Protections, Consumer Reports' Testing Finds. Retrieved from https://www.consumerreports.org/wireless-routers/wireless-routers-lack-basic-security-protections/. Last Accessed: 20 March 2021.
[31] Kjeld Egevang and Paul Francis. 1994. *The IP Network Address Translator (NAT)*. RFC 1631.
[32] Reinaldo Penno et al. 2016. *Updates to Network Address Translation (NAT) Behavioral Requirements*. BCP 127.
[33] Reinaldo Penno et al. 2019. *Operational Security Considerations for IPv6 Networks*. RFC Draft Ver 21.
[34] International Organization for Standardization. 2018. *ISO/IEC: 27000:2018 Information Technology—Security Techniques—Information Security Management Systems—Overview and Vocabulary*. Technical Report.
[35] Bryan Ford, Pyda Srisuresh, and Dan Kegel. 2005. Peer-to-peer communication across network address translators.. In *USENIX Annual Technical Conference, General Track*. 179–192.
[36] UPnP Forum. 2020. UPnP Specification. Retrieved from https://openconnectivity.org/developer/specifications/upnp-resources/.
[37] Marco Ghiglieri, Melanie Volkamer, and Karen Renaud. 2017. Exploring consumers' attitudes of smart TV related privacy risks. In *International Conference on Human Aspects of Information Security, Privacy, and Trust*. Springer, 656–674.
[38] Ryan Giobbi. 2008. UPnP Enabled By Default, SEI Vulnerability Note VU347812. https://www.kb.cert.org/vuls/id/347812/.
[39] Google. 2020. Per-Country IPv6 Adoption. Retrieved from https://www.google.com/intl/en/ipv6/statistics.html#tab=per-country-ipv6-adoption Last Accessed: 20 March 2021.
[40] Google. 2021. Google nest technical specification. Retrieved from https://store.google.com/us/product/nest_wifi_specs. (2021).
[41] Saikat Guha, Kaushik Biswas et al. 2008. NAT Behavioral Requirements for TCP. RFC 5382.
[42] Saikat Guha, Kaushik Biswas et al. 2009. *NAT Behavioral Requirements for ICMP*. BCP 148.
[43] Saikat Guha and Paul Francis. 2005. Characterization and measurement of TCP traversal through NATs and firewalls. In *Proceedings of the 5th ACM SIGCOMM Conference on Internet Measurement*. 18–18.
[44] Fanglu Guo and Tzi-cker Chiueh. 2005. Sequence number-based MAC address spoof detection. In *Proceedings of the International Workshop on Recent Advances in Intrusion Detection*. Springer, 309–329.
[45] Tony Hain. 2000. *Architectural Implications of NAT*. RFC 2993.
[46] Gertjan Halkes and Johan Pouwelse. 2011. UDP NAT and firewall puncturing in the wild. In *Proceedings of the International Conference on Research in Networking*. Springer, 1–12.
[47] Nikolai Hampton and Patryk Szewczyk. 2015. A survey and method for analysing SOHO router firmware currency. https://ro.ecu.edu.au/cgi/viewcontent.cgi?article=1176&context=ism.
[48] Wan Haslina Hassan et al. 2019. Current research on internet of things (IoT) security: A survey. *Comput. Netw.* 148 (2019), 283–294.
[49] Seppo Hätönen, Aki Nyrhinen, Lars Eggert, Stephen Strowes, Pasi Sarolahti, and Markku Kojo. 2010. An experimental study of home gateway characteristics. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*. 260–266.
[50] Ryan Heartfield, George Loukas, Sanja Budimir, Anatolij Bezemskij, Johnny R. J. Fontaine, Avgoustinos Filippoupolitis, and Etienne Roesch. 2018. A taxonomy of cyber-physical threats and impact in the smart home. *Comput. Secur.* 78 (2018), 398–428.
[51] Matt Holdrege and Pyda Srisuresh. 2001. Protocol complications with the IP network address translator (NAT). RFC 3027.

[52] Sebastian Holzapfel, Matthaus Wander, Arno Wacker, and Torben Weis. 2011. SYNI-TCP Hole punching bBased on SYN injection. In *Proceedings of the IEEE 10th International Symposium on Network Computing and Applications*. IEEE, 241–246.
[53] Christopher Huitema. 2006. *Teredo: Tunneling IPv6 over UDP Through Network Address Translations (NATs)*. RFC 4380.
[54] Shiang-Ming Hunag, Quincy Wu et al. 2011. A survey of NAT behavior discovery in VOIP applications. *J. Internet Technol.* 12, 2 (2011), 199–210.
[55] The American Consumer Institute. 2018. *Securing IoT Devices: How Safe Is Your Wi-Fi Router?* Retrieved from https://www.theamericanconsumer.org/wp-content/uploads/2018/09/FINAL-Wi-Fi-Router-Vulnerabilities.pdf.
[56] Senator Hannah-Beth Jackson. 2017. California Senate Bill SB-327, Chapter 866. Retrieved from https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB327. Accessed: 2022-10-17.
[57] Cullen Jennings. 2007. *NAT Classification Test Results*. Internet-Draft draft-jennings-behave-test-results-04. Internet Engineering Task Force. Retrieved from https://datatracker.ietf.org/doc/html/draft-jennings-behave-test-results-04.
[58] Daryl Johnson and Bruce Hartpence. 2010. A re-examination of network address translation security. https://scholarworks.rit.edu/cgi/viewcontent.cgi?article=1764&context=other.
[59] Carolyn Brodie, Clare-marie Karat, John Karat and Jinjuan Feng. 2005. Usable security and privacy: A case study of developing privacy management tools. *ACM International Conference Proceeding Series*, Vol. 93, 35–43. 10.1145/1073001.1073005.
[60] A. Keranen, C. Holmberg, and J. Rosenberg. 2018. *Interactive Connectivity Establishment (ICE)*. RFC 8445.
[61] Kevin S. Killourhy, Roy A. Maxion, and Kymie M. C. Tan. 2004. A defense-centric taxonomy based on attack manifestations. In *Proceedings of the International Conference on Dependable Systems and Networks*. IEEE, 102–111.
[62] Timo Kiravuo, Mikko Sarela, and Jukka Manner. 2013. A survey of ethernet LAN security. *IEEE Commun. Surveys Tutor.* 15, 3 (2013), 1477–1491.
[63] Barbara Kitchenham and Pearl Brereton. 2013. A systematic review of systematic review process research in software engineering. *Info. Softw. Technol.* 55, 12 (2013), 2049–2075.
[64] Deepak Kumar, Kelly Shen, Benton Case, Deepali Garg, Galina Alperovich, Dmitry Kuznetsov, Rajarshi Gupta, and Zakir Durumeric. 2019. All things considered: An analysis of IoT devices on home networks. In *Proceedings of the 28th USENIX Security Symposium*. 1169–1185.
[65] Frank Li and Vern Paxson. 2017. A large-scale empirical study of security patches. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*. 2201–2215.
[66] Xiang Li, Baojun Liu, Xiaofeng Zheng, Haixin Duan, Qi Li, and Youjun Huang. 2021. Fast IPv6 network periphery discovery and security implications. In *Proceedings of the 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'21)*. IEEE, 88–100.
[67] Linksys. 2019. Configuring the MAC Filter feature of the Linksys Smart Wi-Fi Router using the local access interface. Retrieved from https://www.linksys.com/us/support-article?articleNum=143602. Last Accessed: 2 February 2020.
[68] Richard P. Lippmann, David J. Fried et al. 2000. Evaluating intrusion detection systems: The 1998 DARPA off-line intrusion detection evaluation. In *Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX'00)*, Vol. 2. IEEE, 12–26.
[69] R. Mahy, P. Matthews, and J. Rosenberg. 2010. *Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)*. RFC 5766.
[70] Gregor Maier, Fabian Schneider, and Anja Feldmann. 2011. NAT usage in residential broadband networks. In *Proceedings of the International Conference on Passive and Active Network Measurement*. Springer, 32–41.
[71] Paul Marrapese. 2019. IoT Security Flaw Leaves 496 Million Devices Vulnerable At Businesses: Report. Retrieved from https://www.crn.com/news/internet-of-things/300106806/iot-security-flaw-leaves-496-million-devices-vulnerable-at-businesses/-report.htm.
[72] Dylan Martin. 2018. Security Cameras Vulnerable to Hijacking. Retrieved from https://hacked.camera/.
[73] Troy Mattessich. 2012. Exploits and Vulnerabilities of IP Camera's. Retrieved from http://cysecure.org. Accessed: 2021-02-17.
[74] Steve McConnell. 2004. *Code Complete*. Pearson Education.
[75] Jelena Mirkovic and Peter Reiher. 2004. A taxonomy of DDoS attack and DDoS defense mechanisms. *ACM SIGCOMM Comput. Commun. Rev.* 34, 2 (2004), 39–53.
[76] MITRE. 2021. Common Vulnerabilities and Exposures. Retrieved from https://cve.mitre.org/. Last Accessed: 17 February 2021.
[77] H. D. Moore. 2013. Security Flaws in Universal Plug and Play: Unplug. Don't Play. Rapid7. Retrieved from https://information.rapid7.com/rs/411-NAK-970/images/SecurityFlawsUPnP.pdf. Accessed: 2021-02-17.
[78] Philipp Morgner, Christoph Mai, Nicole Koschate-Fischer, Felix Freiling, and Zinaida Benenson. 2019. Security update labels: Establishing economic incentives for security patching of IoT consumer products. Retrieved from https://arXiv:1906.11094.

[79] Andreas Müller, Georg Carle, and Andreas Klenk. 2008. Behavior and classification of NAT devices and implications for NAT traversal. *IEEE Netw.* 22, 5 (2008), 14–19.
[80] Andreas Muller, Nathan Evans, Christian Grothoff, and Samy Kamkar. 2010. Autonomous NAT traversal. In *Proceedings of the IEEE 10th International Conference on Peer-to-Peer Computing (P2P)*. IEEE, 1–4.
[81] Andreas Müller, Andreas Klenk, and Georg Carle. 2008. On the applicability of knowledge based NAT-traversal for home networks. In *Proceedings of the International Conference on Research in Networking*. Springer, 264–275.
[82] Christelle Nader and Elias Bou-Harb. 2021. Revisiting IoT fingerprinting behind a NAT. In *Proceedings of the IEEE International Conference on Parallel and Distributed Processing with Applications, Big Data and Cloud Computing, Sustainable Computing and Communications, Social Computing and Networking (ISPA/BDCloud/SocialCom/SustainCom'21)*. IEEE, 1745–1752.
[83] Netgear. 2019. What is Explicit Beamforming and How Does It Work? Retrieved from https://kb.netgear.com/31299/What-is-explicit-beamforming-and-how-does-it-work.
[84] SMC Networks. 2015. SMC8014WG-SI User Manual. Retrieved from https://manualmachine.com/smcnetworks/ezconnectsmc8014wgsi/479465-user-manual/.
[85] Marcus Niemietz and Jörg Schwenk. 2015. Owning your home nnetwork: Router security revisited. Retrieved from https://arXiv:1506.04112.
[86] NIST. 2021. National Vulnerability Database. Retrieved from https://nvd.nist.gov/. Last Accessed: 17 February 2021.
[87] N. Nthala and Ivan Flechais. 2018. Rethinking home network security. *European Workshop on Usable Security (EuroUSEC'18 23 April 2018, London, England)*. https://dx.doi.org/10.14722/eurousec.2018.23011
[88] National Institute of Standards and Technology. 2022. Recommended criteria for cybersecurity labelling for consumer internet of things (IoT) products. Retrieved from https://doi.org/10.6028/NIST.CSWP.02042022-2. (2022). Last accessed 23 Oct 2022.
[89] Karl Olson, Jack Wampler, Fan Shen, and Nolen Scaife. 2021. NATting else matters: Evaluating IPv6 access control in residential networks. In *Proceedings of the Probability and Meaning Conference (PaM'21)*.
[90] Alexandre Pankratov. 2012. Server-Mediated Setup and Maintenance of Peer-to-Peer Communications. U.S. Patent 8,296,437.
[91] Marc Petit-Huguenin, Gonzalo Salgueiro, Jonathon Rosenberg, Dan Wing, Rohan Mahy, and Phillip Matthews. 2020. *Session Traversal Utilities for NAT (STUN)*. RFC 8489.
[92] Jon Postel. 1981. *Internet Control Message Protocol*. STD 5.
[93] Matt Powell. 2018. *Wi-Fi Router Security Knowledge Gap Putting Devices and Private Data at Risk in UK Homes*. Retrieved from https://www.broadbandgenie.co.uk/blog/20180409-wifi-router-security-survey.
[94] Associated Press. 2006. No Rush to Upgrade Your WiFi Router. Retrieved from https://www.law.com/legaltechnews/almID/1167214009597/?id=1167214009597?id=1167214009597&slreturn=2021101715253. Last accessed 5 Nov 2021.
[95] Philipp Richter, Florian Wohlfart, Narseo Vallina-Rodriguez, Mark Allman, Randy Bush, Anja Feldmann, Christian Kreibich, Nicholas Weaver, and Vern Paxson. 2016. A multi-perspective analysis of carrier-grade NAT deployment. In *Proceedings of the Internet Measurement Conference*. 215–229.
[96] Jonathon Rosenberg. 2010. *Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols*. RFC 5245.
[97] Jonathan Rosenberg, Joel Weinberger, Christian Huitema, and Rohan Mahy. 2003. *STUN—Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)*. RFC 3489.
[98] Tom Rowan. 2010. Negotiating WiFi security. *Netw. Secur.* 2010, 2 (2010), 8–12.
[99] Francisco Ruiz-Aliseda and Peter B. Zemsky. 2006. Adoption is not development: First mover advantages in the diffusion of new technology. INSEAD Business School Research Paper 2007/03 (2006).
[100] Jerome H. Saltzer, David P. Reed, and David D. Clark. 1984. End-to-end arguments in system design. *ACM Trans. Comput. Syst.* 2, 4 (1984), 277–288.
[101] O. Santos. 2016. The Evolution of Scoring Security Vulnerabilities: The Sequel. Retrieved from https://blogs.cisco.com/security/cvssv3-study. Accessed: 2021-02-17.
[102] Rajiv Shah and Christian Sandvig. 2008. Software defaults as de facto regulation the case of the wireless internet. *Info., Commun. Soc.* 11, 1 (2008), 25–46.
[103] Young Sheng, Keren Tan et al. 2008. Detecting 802.11 MAC layer spoofing using received signal strength. In *Proceedings of the IEEE 27th Conference on Computer Communications (INFOCOM'08)*. 1768–1776.
[104] Shodan. 2021. UPnP Exposure Scan. Retrieved from https://www.shodan.io/. Last accessed 20 February 2021.
[105] Hemant Singh, Wes Beebee et al. 2013. *Basic Requirements for IPv6 Customer Edge Routers*. Technical Report. 2070–1721.
[106] Saurabh Singh, Pradip Kumar Sharma, Seo Yeon Moon, and Jong Hyuk Park. 2017. Advanced lightweight encryption algorithms for IoT devices: survey, challenges and solutions. *Journal of Ambient Intelligence and Humanized Computing*, 1–18. https://doi.org/10.1007/s12652-017-0494-4

[107] Nejc Škoberne, Olaf Maennel, Iain Phillips, Randy Bush, Jan Zorz, and Mojca Ciglaric. 2013. IPv4 address sharing mechanism classification and tradeoff analysis. *IEEE/ACM Trans. Netw.* 22, 2 (2013), 391–404.
[108] Matt Smith and Ray Hunt. 2002. Network security using NAT and NAPT. *Proceedings of the 10th IEEE International Conference on Networks (ICON'02): Towards Network Superiority*. 355–360.
[109] Matt Smith and Ray Hunt. 2002. Network security using NAT and NAPT. In *Proceedings of the 10th IEEE International Conference on Networks (ICON'02): Towards Network Superiority*. IEEE, 355–360.
[110] Pyda Srisuresh and Matt Holdrege. 1999. *IP Network Address Translator (NAT) Terminology and Considerations*. RFC 2663.
[111] Günther Starnberger. 2007. *NAT Traversal Techniques in VoIP Protocols*. Ph.D. Dissertation.
[112] Amber Steele. 2016. Keep Your Friends Close and Your Passwords Closer. Retrieved from https://blog.lastpass.com/2016/02/infographic-keep-your-friends-close-your-passwords-closer-2.html/. Last Accessed: 28 March 2022.
[113] Synopsis. 2021. MiniUPnP. Retrieved from https://www.openhub.net/p/miniupnp. Last Accessed: 2021-02-17.
[114] Patryk Szewczyk and Craig Valli. 2009. Insecurity by obscurity: A review of SoHo router literature from a network security perspective. *J. Dig. Forens., Secur. Law* 4, 3 (2009), 1.
[115] Kahkashan Tabassum, Ahmed Ibrahim, and Sahar A. El Rahman. 2019. Security issues and challenges in IoT. In *Proceedings of the International Conference on Computer and Information Sciences (ICCIS'19)*. IEEE, 1–5.
[116] Trendmicro. 2019. UPnP-enabled Home Devices and Vulnerabilities. Retrieved from https://www.trendmicro.com/en_us/research/19/c/upnp-enabled-connected-devices-in-home-upnpatched-known-vulnerabilities.html. Last Accessed: 1 February 2021.
[117] Tripwire. 2014. SOHO Wireless Router (In)Security. Retrieved from http://www.properaccess.com/docs/_SOHO_Router_Insecurity_white_paper.pdf. Last accessed 20 October 2020.
[118] B. Seri and G. Vishnepolsky. 2017. *BlueBorne Technical White Paper*. Armis Labs. Available From: https://info.armis.com/rs/645-PDC-047/images/BlueBorne%20Technical%20White%20Paper_20171130.pdf?_ga=2.119171470.602323090.1679241418-1914181406.1679241418.
[119] Lisandro Ubiedo, Thomas O'Hara, María José Erquiaga, and Sebastian Garcia. 2021. Current state of IPv6 security in IoT. Retrieved from https://arXiv:2105.02710.
[120] Mark Vink, Erik Poll, and Alex Verbiest. 2020. *A Comprehensive Taxonomy of Wi-Fi Attacks*. Ph.D. Dissertation. Radboud University Nijmegen Nijmegen, The Netherlands.
[121] Daniel James Weber. 1998. *A Taxonomy of Computer Intrusions*. Ph.D. Dissertation. Massachusetts Institute of Technology.
[122] P. Weidenbach and J. vom Dorp. 2020. Home Router Security 2020. Retrieved from https://www.fkie.fraunhofer.de/content/dam/fkie/de/documents/HomeRouter/HomeRouterSecurity_2020_Bericht.pdf. Accessed: 2021-02-17.
[123] Christian Wimmer. 2008. *Wireless LAN Security in a SOHO Environment: A Holistic Approach*. GRIN Verlag, 2012.
[124] Dan Wing, Stuart Cheshire et al. 2013. *Port Control Protocol (PCP)*. RFC 6887.
[125] James Woodyatt. 2011. *Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service*. RFC 6092.
[126] Avishai Wool. 2004. A quantitative study of firewall configuration errors. *Computer* 37, 6 (2004), 62–67.
[127] Jie Xiong and Kyle Jamieson. 2013. Securearray: Improving WiFi security with fine-grained physical-layer information. In *Proceedings of the 19th Annual International Conference on Mobile Computing and Networking*. 441–452.
[128] Lixia Zhang. 2008. A retrospective view of network address translation. *IEEE Netw.* 22, 5 (2008), 8–12.
[129] Lizhuo Zhang, Weijia Jia et al. 2010. Research of TCP NAT traversal solution based on port correlation analysis & prediction algorithm. In *Proceedings of the 6th International Conference on Wireless Communications Networking and Mobile Computing (WiCOM'10)*. IEEE, 1–4.

Received 1 April 2022; revised 19 November 2022; accepted 21 February 2023