Plight at the End of the Tunnel
Legacy IPv6 Transition Mechanisms in the Wild

John Kristoff, Mohammad Ghasemisharif, Chris Kanich, and Jason Polakis

University of Illinois at Chicago
{jkrist3,mghas2,ckanich,polakis}@uic.edu

Abstract. IPv6 automatic transition mechanisms such as 6to4 and ISATAP endure on a surprising number of Internet hosts. These mechanisms lie in hibernation awaiting someone or something to rouse them awake. In this paper we measure the prevalence and persistence of legacy IPv6 automatic transition mechanisms, together with an evaluation of the potential threat they pose. We begin with a series of DNS-based experiments and analyses including the registration of available domain names, and demonstrate how attackers can conduct man-in-the-middle attacks against all IPv6 traffic for a significant number of end systems. To validate another form of traffic hijacking we then announce a control set of special-purpose IPv6 prefixes that cannot be protected by the RPKI to see these routes go undetected, accepted, and installed in the BGP tables of over 30 other upstream networks. Finally, we survey the Internet IPv4 address space to discover over 1.5 million addresses are open IPv6 tunnel relays in the wild that can be abused to facilitate a variety of unwanted activity such as IPv6 address spoofing attacks. We demonstrate how many attacks can be conducted remotely, anonymously, and without warning by adversaries. Behind the scenes our responsible disclosure has spearheaded network vendor software updates, ISP remediation efforts, and the deployment of new security threat monitoring services.

1 Introduction

The meteoric rise of the Internet motivated the proposal of IPv6 over two decades ago. However, rather than decree an instantaneous conversion and face the unavoidable disruption that would ensue, a slow migration started around the turn of the century and is still underway [33]. Around 25 years later, reports from Akamai [7] and Google [26] suggest that over 25% of client systems are using IPv6 in 2020. While IPv6 adoption has been substantial, a significant majority of users lack IPv6 connectivity. The slow migration necessitated the design, implementation, and deployment of transition mechanisms in order to maintain reachability between communicating hosts that lack direct connectivity to each other using their chosen version of IP.

A handful of security-related concerns about transition mechanisms were documented in IETF RFCs after the technology first arose [34, 39]. These early concerns mentioned the lack of authentication on endpoints, and how they can be

used for launching distributed denial-of-service attacks or IPv4 policy avoidance. The referenced RFCs summarize certain potential security threats, but do not provide specific guidance on how to mitigate them. Despite the concerns, these mechanisms were still added to commodity operating systems. More importantly, reports had not envisioned the DNS-based attacks and the extent of implementation weaknesses we uncover in this work. Based on reports [42, 44] that provide statistics on the versions of users' Windows operating system, $\sim 33\%$ *of all Windows machines in the wild* currently have these automatic transition mechanisms *enabled by default*. Transition mechanisms are also supported by almost all other major operating systems.

Since transition mechanisms are intended to work around the shortcomings of a local IPv4-only network connection, many of them were designed as host-initiated tunneling protocols. Tunnels are the most straightforward solution to the network protocol transition problem, but as we will show, the implementation of the IPv6 automatic transition mechanisms were designed with little consideration for long term effects of on-by-default settings or the ease at which man-in-the-middle (MitM) attacks can be conducted.

We present several techniques that allow a remote attacker to meet the preconditions for activating the transition mechanism implementations undetected. Activation enables attackers to perpetrate stealthy traffic hijacking on a significant population of Windows hosts where these mechanisms currently lie dormant. Furthermore, with IPv6 being the generally preferred network layer protocol when given the choice between IPv6 and IPv4 [41], transition mechanism tunnels will handle a large portion of the network traffic for hosts without native IPv6 connectivity. The attack's impact may be further amplified by manipulating unauthenticated DNS responses that traverse a malicious tunnel by adding or including AAAA answers, thus "guiding" the client towards more IPv6 destinations.

In this paper, we investigate the persistence of *legacy* transition mechanisms by conducting a series of measurements, including a longitudinal study over the course of 13 months using data from multiple network vantage points, detailing the severity and feasibility of different attacks.

We consider two different attack vectors that capture adversaries with vastly different capabilities and resources. First, we demonstrate how an attack that requires only a domain name registration allows an adversary to hijack the IPv6 traffic for a substantial number of hosts having a domain suffix in a zone we control or can register a name for. We can directly observe 32,156 hosts susceptible to IPv6 traffic hijacking using this technique. The only additional requirement for this attack is the absence of network address translation (NAT). While these vulnerabilities can be directly exploited through the registration of specific domain names, we also explore a more sophisticated attack, where an adversary can announce BGP routes into the Internet routing tables.

Lastly, we conduct an Internet survey of open relay tunnels in the wild and how they can be used as a springboard for attacks, such as traffic reflection, spoofing, or the discovery of private network infrastructure. Our Internet-wide

scans reveal over 1.5 million IPv4 addresses that can be exploited for such attacks, with further investigation revealing a portion of those relays consist of a widely deployed backbone router that allows IPv6 tunneling for anyone on the Internet by natively forwarding encapsulated IPv6 traffic arriving on an IPv4 interface. To the best of our knowledge, we are the first to report on these IPv4 hosts functioning as open IPv6 tunnel relays and the potential for misuse. Overall, our study sheds light on new attack vectors that pose a significant threat to the Internet, and highlights the importance of mobilizing the networking and operational security communities for deploying appropriate countermeasures.

In summary, our research contributions are as follows:

- We conduct a measurement of the contemporary use of legacy IPv6 automatic transition mechanisms within end hosts, transition mechanism-providing servers, and network infrastructure. We find multiple MitM attack vectors which are enabled by default on millions of Internet-connected hosts including DNS-based vectors unanticipated by the original designers or earlier reports.
- We further explore the practical implications of these MitM attacks, both in scope and severity. Our experimental analysis, driven by data collected from academic institutions, ISPs, and other organizations reveals the magnitude of the threat.
- Due to the severity of these vulnerabilities, we have reported them and coordinated with various trusted communities of network administrators of vulnerable networks, router vendors, and multiple incident response and threat intelligence reporting organizations. We also discuss countermeasures and mitigation strategies.

2 Background

A full accounting of all IPv6 transition mechanisms is beyond the scope of this paper. For example, newer mechanisms such as 6rd [45], DS-Lite [21], and 464XLAT [36] are not considered here. Instead, we focus on a subset of *legacy* automatic transition mechanisms. Three of the earliest and most popular are 6to4, ISATAP, and Teredo. Their peculiar use of specific address prefixes, the DNS, or tunnel bootstrapping allows us to conduct extensive measurements and experiments demonstrating their susceptibility to various forms of attack.

6to4. IETF RFC 3056 [9] describes one of the earliest automatic transition mechanisms for IPv6 in IPv4 tunneling. The Internet Assigned Numbers Authority (IANA) designated the `2002::/16` prefix to be used by 6to4 systems. [30] Bits 17 to 48 of a 6to4 address correspond to the globally unique 32-bit IPv4 address of the 6to4 host or site network. Systems behind a network address translator (NAT) or using private addresses cannot use 6to4. A 6to4 system can communicate with IPv6 over an intermediate IPv4 subpath by encapsulating IPv6 packets in IPv4 towards a well-known destination address from the IANA-designated special-use anycast prefix `192.88.99.0/24`. Any network announcing this prefix must be willing to accept a 6to4 system's encapsulated packet, re-

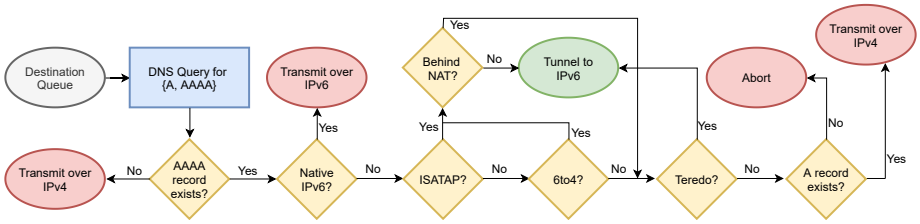

Fig. 1: Conceptual flowchart outlining the conditions for an IPv6 transition mechanism to be used.

move the outer IPv4 layer, and relay the enclosed IPv6 packet towards the IPv6 destination directly. Conversely, traffic back to a 6to4 system needs a relay that can add the IPv4 encapsulation onto an IPv6 datagram so that it may continue on the subpath that is IPv4-only. Like the IPv4 anycast prefix, a network advertising 2002::/16 must be willing to perform this relay service in the opposite direction.

ISATAP. Where 6to4 relies on the global routing infrastructure with well known prefixes and addressing for host configuration and packet forwarding, the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) is widely implemented with the help of the DNS to automatically construct an IPv6 over IPv4 tunnel [25]. A typical ISATAP client issues a DNS `A query` for a name with the suffix of the locally defined zone and a label prefix of *isatap* (e.g., `isatap.myzone.example.net`). If an address is indeed returned for this name, an ICMPv6 router solicitation and ICMPv6 router advertisement are exchanged over an IPv4 tunnel. The ISATAP client uses the source of the router advertisement response as the default IPv6 tunnel relay router. As with 6to4, ISATAP only works on hosts that are not behind a NAT.

Teredo. IETF RFC 4380 [29], describes an IPv6 over UDP-based automatic transition mechanism intended for clients behind a NAT device. Teredo clients communicate with a configured or discovered tunnel relay like ISATAP, except they do not need a public IPv4 address. Teredo client IPv6 addresses are derived from a combination of server and client attributes appended to the well-known Teredo prefix (2001::/32).

Interface and mechanism selection. It is possible for a system to have multiple active IPv6 interfaces and addresses. When such an IPv6 host has traffic to deliver, it must select an interface and source address from which to send traffic. IETF RFC 6724 [43] outlines the default address selection algorithms that should be used. If multiple transition mechanisms are active and available on a host without native IPv6 connectivity, traffic delivery and reception will tend to use only one available transition mechanism at a time. To illustrate interface selection, Figure 1 depicts how a Microsoft Windows system will make it's choice. As shown in the top left, a client commonly starts communication with a DNS query. If an IPv6 address answer is provided, the most preferred interface type that is available will be used, falling back to IPv4 as a last resort.

Network connectivity status indicator. When joining new networks or activating a new network interface, Microsoft Windows machines perform a connectivity test with a HTTP GET request for `www.msftncsi.com/ncsi.txt`. If this test succeeds the interface is assumed functional for as long as the interface state remains active. In order to more accurately measure IPv6 transition mechanism usage in our tunnel relay experiments, we want to ensure this test is successful. Microsoft systems will always attempt to use an IPv6 interface that passed the "ncsi" test, but will fall back to IPv4 with little to no perceived interruption of service if IPv6 communications fail.

3 Methodology & Data

In this section we provide a high-level overview of our experimental setup, methodology, limitations and network vantage points for measuring transition mechanism behavior, security threats, and privacy risks.

Domain Registrations. We registered dozens of *isatap* names in EDU-A, EDU-B, top-level domains (TLDs) and shared domain providers. Where possible we ran our own authoritative name servers for these names with the EDNS0 client subnet option [13] and extensive logging configured. These vantage points provide a diverse, but limited view of the global DNS name space and client population for our experiments.

EDU-A Functional ISATAP Relay. We setup a fully functional ISATAP relay that handed out public IPv6 /64 prefixes and relayed tunneled traffic received from any of the institution's client population that had ISATAP enabled by default.

EDU-B Dysfunctional ISATAP Relay. This relay was configured to receive tunnel requests for all clients within the institution's primary DNS domain and the computer science domain. It was also the relay for a sample of ISATAP domains we registered in a number of TLDs and dynamic DNS providers. This tunnel relay system operated in the "dysfunctional" state, which would appear as a valid IPv6 path, but would ultimately reject traffic not associated with tunnel establishment and control.

DNS Query Logs. From EDU-A we received anonymized client query logs from their local resolvers. A large U.S. cable modem operator also provided us with anonymized DNS query log data for a large city service area containing any IPv6 transition mechanism label in the query name. We also leverage the DNS query data collected by the DNS-OARC DITL project for the two prior two years available. [20]

BGP Route Announcements. We coordinated with the EDU-A upstream research and education network (REN) to announce five distinct IPv6 subprefixes in the `2002::/16` 6to4 block corresponding to three EDU-A IPv4 prefixes (a /16, /18, and /20). This allowed us to measure the potential to launch a global IPv6 transition mechanism hijack without altering the path of any actual traffic.

Internet-wide Scans. We survey the entire Internet IPv4 address space for open and accessible IPv6 tunnel relay services. We first issue a single ICMPv6

router solicitation encapsulated in IPv4 to discover any open ISATAP relays. We then issue an ICMPv6 echo request encapsulated in IPv4. These ICMPv6 messages uncover either 6to4 or raw protocol 41 processing nodes if we receive a corresponding ICMPv6 echo response at our control IPv6 destination.

Ethical and Privacy Considerations. Our experiments required careful planning and review to steer clear of compromising user privacy and to avoid adversely altering Internet application functionality. We performed several internal experiments to ensure that global experiments would not negatively impact users' connectivity or privacy. In all but the experiments being led and controlled by EDU-A, our experiments are limited to tunnel discovery and bootstrapping traffic. To ensure that there is no violation of the privacy of users, all data collection scripts aggregated and anonymized the results (raw data was not retained) without human intervention and the data collection was operated by computing support staff who verified the code's operation and only provided the anonymized, aggregated results to the research team. EDU-A deployed a local, production ISATAP relay, to which we had no direct access, in order to establish ground truth and ensure our attack scenarios could be carried out on real application traffic in practice without user intervention. We submitted a detailed description of our experimental protocol to our university's IRB prior to any experiments, and they determined that this research does not qualify as incorporating human subjects.

4 Analysis

The various legacy automatic transition mechanisms we examine are architecturally similar. They each consist of two fundamental types of systems, tunnel clients and tunnel relays. We present our analysis by examining the threats from the perspective of each system type. The primary vulnerability tunnel clients face is the threat of stealthy man-in-the-middle attacks on all traffic bound for IPv6 hosts. Tunnel relays on the other hand can be impersonated and abused. Impersonation attacks against relays can be enable MitM attacks against tunnel clients. Moreover, tunnel relay abuse can enable various kinds of unwanted activities such as service theft or origination spoofing attacks.

4.1 Attacks Against Tunnel Clients

DNS Capture. Since ISATAP clients typically perform a look up based on the client's default zone, we focus our attention on this mechanism where an attacker could most easily gain access to a number of zones without raising suspicion. Other types of DNS capture attacks, such as cache poisoning could also be used.

Our registered ISATAP domain names received approximately three million queries per month between April 2018 and May 2019. Recall that a DNS query is the first step in bootstrapping an ISATAP tunnel client. Until we registered

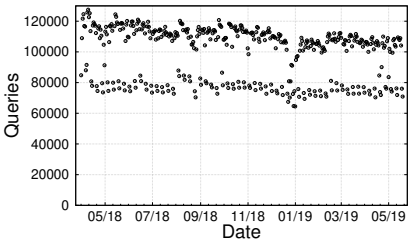

Fig. 2: Name queries received by our authoritative DNS resolvers for our registered ISATAP names, for one year.

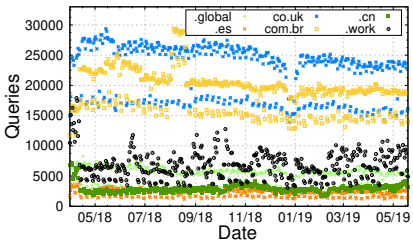

Fig. 3: Most popular ISATAP name queries received by our authoritative DNS resolvers, for one year.

these names, the queries likely went unanswered and the client's ISATAP bootstrapping process ceased until a change in interface state restarted the process. The daily volume shown in Figure 2 exhibits a noticeable work week oscillation, suggesting that many of these queries originate from end user systems that tend to go offline during the weekend. We observe a slow decline towards the end of the monitoring period, which may correspond with the roll out of new systems that have ISATAP disabled by default.

We also break down the queries for the six most popular domain names we registered in Figure 3. The fact that these domains receive thousands of ISATAP queries per day suggests the relative frequency for more popular domains (e.g., ending in `.comcast.net`) will likely be orders-of-magnitude higher. We believe that our top level names are a relatively small sample of the coverage that attacks leveraging IPv6 automatic transition mechanisms can reach. When we examined the DNS-OARC DITL data we see relatively few ISATAP queries for existing zones, but tens of millions of queries for vendor, special-use, or names in private domain over the course of just two days.

In our experiments almost 3K out of 163K resolvers supply EDNS0 client subnet option data. While that is a small fraction of the total number of resolvers observed, over 30% of all queries contain client subnet data. This is due to the disproportionate query volume Google contributes, as their resolvers supply the client subnet data by default if they detect it is supported at the authoritative server. The distinct number of client subnets we see over the course of one year's worth of queries is 96,061. We geo-locate these client subnets to their country of origin and find that they are located all over the world. The extensive use of third-party DNS resolvers (e.g., Google) highlights that these entities are well positioned to impose protections.

We also examine transition mechanism queries seen at EDU-A and a cable modem ISP's resolvers by Microsoft Windows clients for one day in Figure 4. This includes type A (IPv4 address mapping) queries for any name with the *isatap* prefix label, and the fully qualified domain names *6to4.ipv6.microsoft.com* or *teredo.ipv6.microsoft.com*. The cable modem ISP client population is largely be-

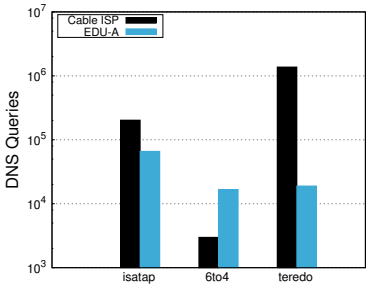

Fig. 4: Comparison of DNS queries for two networks.

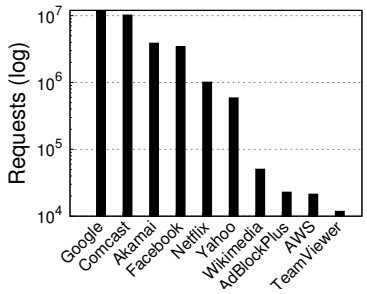

Fig. 5: Top 10 destination domain requests seen at EDU-A.

hind consumer-grade NAT devices. This is reflected in the proportionally higher number of Teredo queries seen at the ISP than EDU-A. Nevertheless, ISATAP queries still make up a large amount of the transition mechanism activity observed in both environments.

Relay Capture. We extend our analysis of threats against clients with the operation of the EDU-A Functional and EDU-B Dysfunctional ISATAP tunnel relays. Figure 5 summarizes the most popular network destinations EDU-A ISATAP tunnel client users were destined for over the course of one 24-hour period. The EDU-A network operations team reviewed these traffic patterns and they believed them to be expected client system traffic behavior that was running over IPv6 instead of IPv4. The traffic includes various forms of email communications, social networking, e-commerce activity, and scientific research.

Client connections to the EDU-B ISATAP relay came primarily, but not entirely, from the EDU-B user population. The ISATAP names registered in co.uk and net.br were also a popular source of ISATAP clients. The most popular IPv6 destinations from clients were concentrated at popular web hosting properties such as Google, Cloudflare, Microsoft, and a handful of content distribution providers. While most attempted traffic through the relay can be classified as HTTP(S), we also observed DNS, FTP, NTP, SMTP, SNMP, SSH, and VPN tunnel attempts. Figure 6 shows the eight destination ports that received the highest number of client connections from a random representative weekday, and that a significant amount of unencrypted HTTP and DNS traffic would be visible to a hypothetical attacker. Since most DNS stub resolvers do not perform DNSSEC-based authentication of answers, attackers could filter out A responses and leave only the valid AAAA answers, forcing all victim traffic to IPv6 capable hosts to transit the malicious tunnel.

Route Hijacking. In April 2019 we began originating five more-specific BGP routes within the 6to4 `2002::/16` prefix from EDU-A. We wanted to evaluate whether we could successfully conduct a targeted attack against 6to4 traffic. The upstream REN agreed to allow these prefixes into their backbone, but they limited the propagation to a subset of regional REN participants for safety rea-

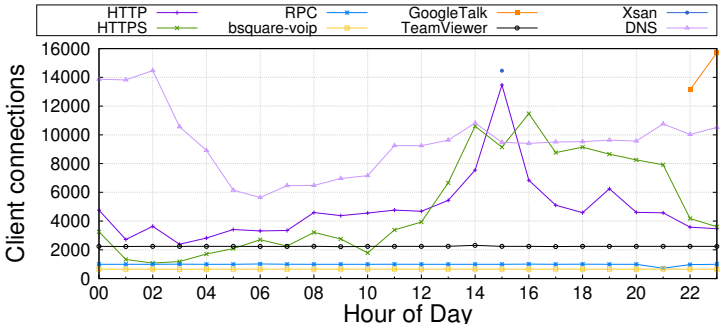

Fig. 6: Top destination services based on port numbers.

sons; they were not relayed to commercial or international peering partners. Nonetheless, at least twelve REN participant networks installed these routes into their routing tables, and RouteViews [3] observed the route from over 30 networks, including multiple tier 1 ISPs. This experiment demonstrates that while customer route filtering may be common practice for ISPs, route filtering between large ISPs and RENs is typically less strict and often inconsistent. Since the IPv6 transition prefixes are not currently protected by the RPKI [35] and may be announced by any origin network, the feasibility of traffic capture is even easier than traditional unicast route hijacking. This experiment ran for many months and to the best of our knowledge there were no public reports or inquiries about the nature of these spurious announcements.

4.2 Attacks Against Tunnel Relays

Theft of Service. In November 2018 and April 2019 we surveyed the entire IPv4 address space for ISATAP-compatible open tunnel relays on the public Internet. The 2018 survey recorded 765 ICMPv6 router advertisement (RA) responses while the 2019 survey recorded 628 responses, totaling 841 unique addresses. Further examination suggests that these hosts are mostly Microsoft IIS web servers with firewalls disabled and forwarding capability for remote hosts activated. Their router advertisement responses typically include a number of available routes, most commonly 6to4 prefixes, but also some Teredo and unique local IPv6 unicast prefixes [27]. We did not find any probed hosts offering unique global IPv6 addresses. We classify these as ISATAP-capable since our client was able to successfully self-configure using the ISATAP address acquisition process.

Upon seeing how most open ISATAP tunnel relays would provide 6to4 addresses by default, in April 2019 we issued ICMPv6 echo request messages, encapsulated in IPv4 protocol 41 packets, to the entire IPv4 address space. Much to our chagrin we discovered 1,546,843 IPv4 addresses around the globe would relay the enclosed IPv6 message to the intended destination. We were surprised that such a large number of system configurations not only had the 6to4 mechanism enabled, but were left unprotected on the public Internet, allowing anyone

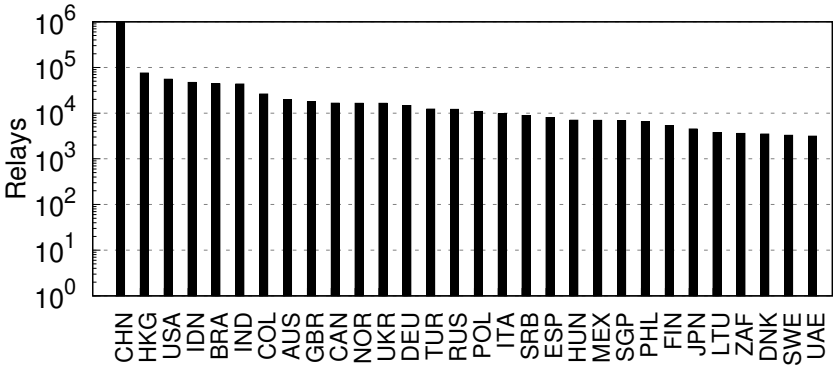

Fig. 7: Open 6to4 relays' country of origin.

to relay traffic through them. We break down those relays according to their geographic distribution and find that China, Hong Kong, USA, Indonesia and Brazil have the most relays. We provide a list of the top 30 countries in Figure 7.

We ran an *nmap*[24] survey on a sample of these open IPv6 tunnel relays and discovered an alarming number of fingerprints matched backbone routers from one of the world's largest networking vendors, which we confirmed through two different network operators. We estimated that approximately 7% of the addresses discovered were from this particular vendor. This particular brand of equipment exhibited particularly curious behavior. They process IPv4 protocol 41 datagrams by first removing the IPv4 header. Then the IPv6 destination address is consulted and the IPv6 datagram is forwarded along its way. In other words, this particular class of equipment acts as an IPv6 default router for any IPv6 traffic it receives, even if encapsulated in IPv4 first. This led to additional interesting observations, two of which we briefly describe below.

Infrastructure Abuse. These vendor backbone routers had a noteworthy peculiarity. They were rarely configured to support the 6to4 mechanism. Therefore, if the embedded IPv6 destination is the 6to4 equivalent destination of the backbone router's own IPv4 address, the packet will attempt to follow whatever path the router has to the `2002::/16` prefix. A 6to4 network service provider upon receiving this packet will examine the IPv6 destination, put the IPv6 message into an IPv4 packet and send it back towards the router's IPv4 address where the process repeats until the enclosed IPv6 hop limit field eventually expires, but not before the packet iterates through this loop. This leads to a potential DoS attack where sending one tunnel packet can expand to multiple packets cycling in a loop between the 6to4 gateway and the backbone router.

Infrastructure Disclosure. Another observation from those backbone routers appeared when we attempted to evaluate the IPv6 path of the aforementioned loops. Output from traceroute often showed our relayed packets were able to traverse IPv6 paths not accessible via the native public IPv6 Internet. Listing 1.1 shows the partial path through an IPv6 open relay on a North American ISP

```
$ traceroute -n -q1 2002:c000:0201::1
traceroute to 2002:c000:0201::1,
30 hops max, 80 byte packets
1  *
2  2001:4958:300:449::b  63.779 ms
3  2001:4958:300:449::a  63.764 ms
4  ::ffff:64.230.193.173  70.472 ms
5  *
6  2001:4958:300:d::1b  64.748 ms
7  2001:470:1:802::1  64.122 ms
```

Listing 1.1: Traceroute traversing hidden paths.

network (the target destination address has been anonymized). In this example, traceroute should have terminated at the first hop. However, this class of equipment blindly forwards this packet along a path towards a route advertised for the `2002::/16` prefix. The fourth hop shows an IPv4-mapped address, which should not appear on the public Internet. Access to addresses and paths intended for internal-only use may facilitate network reconnaissance or attacks that bypass security policies.

Origination Spoofing. Open tunnel relay systems not only allowed us to obtain an IPv6 address and relay traffic through them, they facilitated source IP address spoofing. Most operators of networks, where directly attached hosts emit packets, perform a form of source address validation (SAV) on IP datagrams at the first hop ingress router. [5] However, routers only perform this validation on the outer IP layer, not on the IPv6 source address of encapsulated packets. We were able to set the IPv6 source address to most any value of our choosing. The tunnel relays re-encapsulate our original IPv6 datagrams inside a new outer IPv4 header using the tunnel's IPv4 source address before relaying further. By the time a 6to4 gateway finally receives the packet, all SAV of the IPv6 address has been bypassed. If coupled with a reflection and amplification style attack, this behavior can significantly complicate denial-of-service attack mitigation.

5 Discussion

Susceptible Population. According to online reports, over 30% of Microsoft Windows machines in the wild [42] run OS versions up to Windows 8.1, which have these mechanisms enabled *by default*. While more recent versions of Windows have begun to disable all legacy IPv6 transition mechanisms, the functionality still remains in the operating system. Judging by the significant volume of DNS queries we see for ISATAP names and the vast number of open tunnel relays on many other types of systems, we can safely conclude these mechanisms stubbornly persist, posing risks not only to the systems and users, but to the entire Internet.

Countermeasures. Effective mitigation strategies require significant global coordination as we outline below. We summarize various mitigations that can be implemented to prevent exploitation of these transition mechanisms.

Protocol 41 is used to identify whether an IPv6 datagram is encapsulated within an IP payload, and is at the epicenter of these transition mechanisms. Limiting the transmission of protocol 41 packets would mitigate most attacks we uncover. Some part of transition mechanism bootstrapping, such as DNS queries, may continue unfettered, but would be rendered largely ineffective if protocol 41 packets cannot be relayed.

DNS. As we have shown, the most popular laptop and desktop operating system has made extensive use of DNS to locate and prepare IPv6 links. This feature is susceptible to attacks from off-path attackers. However, the DNS infrastructure is also a place to apply control and policies. Operators of resolvers can exert control over these well known transition mechanism names, either by configuring a local authoritative zone for the names or using response policy zones (RPZ) [47], to render them inactive. Domain name registries, registrars, and ICANN could institute policies to declare certain special-use names as off-limits for registration.

Routing. Legacy IPv6 automatic transition mechanisms such as 6to4 and Teredo utilize well known address prefixes. The routing system provides an operationally centralized means of control to monitor and limit the dissemination of route announcements covering the well known transition address space.

OS and Network Configurations. It is a positive step that Microsoft has disabled these mechanisms in recent versions of their OS. However, reports of older Windows hosts in the wild and our measurements indicate that millions of systems still have these mechanisms turned on. Furthermore, the vast number of open tunnel relays we identified are rarely Windows systems, highlighting the fact that automatic transition threats span a variety of operating systems and device types. These automatic mechanisms can be disabled (or removed) from individual systems by default.

Responsible Disclosure. We have proactively engaged the vendor and operational community for mitigating the attacks we described that can target publicly vulnerable systems. After months of verification, software refactoring, and testing the routing operating system code, a router vendor issued a "high" alert encouraging customers to upgrade or apply configuration work-arounds to avoid the vulnerability. Another hardware vendor verified an issue with their equipment and sent us one to further evaluate in our lab. We also leveraged our personal contacts in the incident response and network security community [1, 2, 23, 31] to coordinate the responsible disclosure of our findings to other vendors and operators affected by the suite of threats we uncovered. Our findings have also renewed discussions in the IPv6 community to officially deprecate these transition mechanisms, encouraging their removal not only from service, but from being made available in systems even when not enabled by default. One of the largest 6to4 service providers has also informed us they are considering a complete shut down of their relay service. Finally, with the help of threat intel-

ligence reporting organizations [4, 6, 14], notifications for systems identified to be at risk can be disseminated to administrative contacts before these findings enter the public sphere. These organizations can also use our findings to build automated scanning and alerting reports for the Internet community at large.

6 Related Work

IPv6 concerns. Ullrich et al. [46] provide a broad overview of security and privacy concerns related to IPv6, and while they mention tunneling between IP protocols, they do not mention the lack of authentication on tunnel creation that enables the attacks we describe.

IPv6 as an evasion technique. Carter [11] warned of attackers setting up proxy interfaces to relay traffic between IPv4 and IPv6 hosts. US-Cert [18] drew attention to malware that enables IPv6 transport, including automatic transition mechanisms, to evade IPv4-only defenses. Blumbergs et al.[8] discussed the limitations intrusion detection systems have when IPv6 transition mechanism tunnels are used for data exfiltration. Czyz et al. [17] highlighted the discrepancies in the access to specific ports. Hong et al. [28] found several vulnerabilities in cellular networks.

Measuring of transition mechanisms. While IPv6 deployment has increased in recent years [12, 15, 19], the underlying factors influencing its adoption [38] indicate that it's unlikely that IPv4 will disappear anytime soon. Czyz et al.[16] deployed an IPv6 sensor on unused address space to observe unsolicited activity. In a similar study, Karis et al.[32] conducted active measurements of IPv6-enabled web clients. Elrich et al.[22] explored the behavior and traffic patterns seen by active Teredo and 6to4 clients on a large academic network and compared them to automatic tunneling mechanisms and native IPv4 communications. Savola compiled a number of observations in the operation of a large, public 6to4 relay service [40] and characterized client system behavior and traffic patterns.

Traffic hijacking attacks. Very similar to the MitM attacks we describe are the Chen et al. [10] hijacking attacks enabled by the Web Proxy Auto-Discovery (WPAD) protocol. Nakibly and Arov discussed a class of routing loop attacks using IPv6 tunnels [37], which took advantage of inconsistency between different transition technologies.

7 Conclusion

We presented a comprehensive exploration of legacy IPv6 transition mechanisms on the Internet along with a series of experiments demonstrating the security and privacy risks they continue to pose. We conducted a study using data collected from multiple network vantage points and found a significant number of hosts run operating systems with IPv6 automatic transition mechanisms enabled by default. These mechanisms often lie dormant, idling by until the right set of circumstances triggers their use. If an attacker provisions the necessary resources

or successfully positions themselves in the network path, they can covertly intercept all IPv6 traffic, including traffic towards critical and high-value services like Google and Facebook. Our DNS registration and route announcement experiments explored the practicality and feasibility of different attack vectors that capture adversaries of varying sophistication and resourcefulness. Furthermore, we found a significant number of open tunnel relays, including many on high-cost specialized ISP backbone routers that can facilitate a wide range of attacks such as IPv4 address spoofing and policy bypass. While we have set things in motion by disclosing our findings to certain network administrators, hardware vendors, ISPs, and incident reporting organizations, we hope to bring more attention to the prevalence and risk of legacy IPv6 automatic transition mechanisms in order to accelerate their extinction and countermeasures.

Acknowledgments

We would like to thank our shepherd Ioana Alexandrina Livadariu and the anonymous reviewers for their valuable feedback. We are grateful and indebted to a number groups and individuals that helped make this work possible. The ACCC at UIC, IS at DePaul University, DNS-OARC, Randy Bush, Brian Carpenter, Geoff Huston, Alex Latzko, and the anonymous network operators and vendors who helped validate our findings, or worked to mitigate potential problems. This work was partially supported by the National Science Foundation under contract CNS-1934597. Any opinions, findings, conclusions, or recommendations expressed herein are those of the authors, and do not necessarily reflect those of the US Government.

Bibliography

[1] NSP Security Forum, https://puck.nether.net/mailman/listinfo/nsp-security

[2] Ops-Trust, https://portal.ops-trust.net

[3] RouteViews, https://www.routeviews.org

[4] Shadowserver, https://www.shadowserver.org/

[5] Spoofer: Protect your network and the global internet, https://spoofer.caida.org

[6] Team Cymru, https://www.team-cymru.com/

[7] Akamai: Akamai State of the Internet IPv6 Visual Adoption, https://www.akamai.com/us/en/resources/our-thinking/state-of-the-internet-report/state-of-the-internet-ipv6-adoption-visualization.jsp

[8] Blumbergs, B., Pihelgas, M., Kont, M., Maennel, O., Vaarandi, R.: Creating and Detecting IPv6 Transition Mechanism-based Information Exfiltration Covert Channels. In: In Proceedings of the Nordic Conference on Secure IT Systems. pp. 85–100. NordSec 2016 (2016)

[9] Carpenter, B.E., Moore, K.: Connection of IPv6 Domains via IPv4 Clouds. RFC 3056, RFC Editor (February 2001), https://rfc-editor.org/rfc/rfc3056.txt

[10] Chen, Q.A., Osterweil, E., Thomas, M., Mao, Z.M.: MitM Attack by Name Collision: Cause Analysis and Vulnerability Assessment in the New gTLD Era. In: In Proceedings of the 2016 IEEE Symposium on Security and Privacy. pp. 675–690. S&P 2016 (2016)

[11] Cisco: Securing IPv6 Transition Technologies (2011), https://web.archive.org/web/http://blogs.cisco.com/security/securing-ipv6-transition-technologies

[12] Colitti, L., Gunderson, S.H., Kline, E., Refice, T.: Evaluating IPv6 Adoption in the Internet. In: International Conference on Passive and Active Network Measurement. pp. 141–150. PAM 2010 (2010)

[13] Contavalli, C., Gaast, W.v.d., Lawrence, D.C., Kumari, W.: Client Subnet in DNS Queries. RFC 7871, RFC Editor (May 2016), https://rfc-editor.org/rfc/rfc7871.txt

[14] CyberGreen: Cybergreen, https://www.cybergreen.net

[15] Czyz, J., Allman, M., Zhang, J., Iekel-Johnson, S., Osterweil, E., Bailey, M.: Measuring IPv6 Adoption. In: In Proceedings of the SIGCOMM Conference. pp. 87–98. SIGCOMM 2014 (2014)

[16] Czyz, J., Lady, K., Miller, S.G., Bailey, M., Kallitsis, M., Karir, M.: Understanding IPv6 Internet Background Radiation. In: In Proceedings of the 2013 Internet Measurement Conference. pp. 105–118. IMC 2013 (2013)

[17] Czyz, J., Luckie, M., Allman, M., Bailey, M.: Don't Forget to Lock the Back Door! A Characterization of IPv6 Network Security Policy. In: In

Proceedings of the Network and Distributed System Security Symposium. NDSS 2016 (2016)
[18] Department of Homeland Security: Malware Tunneling in IPv6 (2012), https://www.us-cert.gov/security-publications/malware-tunneling-ipv6
[19] Dhamdhere, A., Luckie, M., Huffaker, B., Claffy, K., Elmokashfi, A., Aben, E.: Measuring the Deployment of IPv6: Topology, Routing and Performance. In: In Proceedings of the 2012 Internet Measurement Conference. pp. 537–550. IMC 2012 (2012)
[20] DNS-OARC: DITL Traces and Analysis, https://www.dns-oarc.net/oarc/data/ditl
[21] Durand, A., Droms, R., Woodyatt, J., Lee, Y.L.: Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion. RFC 6333, RFC Editor (August 2011), https://rfc-editor.org/rfc/rfc6333.txt
[22] Elich, M., Velan, P., Jirsik, T., Celeda, P.: An Investigation Into Teredo and 6to4 Transition Mechanisms: Traffic Analysis. In: In Proceedings of the 7th IEEE Workshop on Network Measurements. pp. 1018–1024. WLN 2013 (2013)
[23] FIRST.org: Forum of Incident Response Security Teams, https://www.first.org
[24] Fyodor: Nmap - network mapper, https://nmap.org
[25] Gleeson, T., Thaler, D., Templin, F.: Intra-Site Automatic Tunnel Addressing Protocol (ISATAP). RFC 5214, RFC Editor (March 2008), https://rfc-editor.org/rfc/rfc5214.txt
[26] Google: Google IPv6 Statistics, https://www.google.com/intl/en/ipv6/statistics.html
[27] Hinden, R.M., Haberman, B.: Unique Local IPv6 Unicast Addresses. RFC 4193, RFC Editor (October 2005), https://rfc-editor.org/rfc/rfc4193.txt
[28] Hong, H., Choi, H., Kim, D., Kim, H., Hong, B., Noh, J., Kim, Y.: When Cellular Networks Met IPv6: Security Problems of Middleboxes in IPv6 Cellular Networks. In: In Proceedings of the IEEE European Symposium on Security and Privacy. pp. 595–609. Euro S&P 2017 (2017)
[29] Huitema, C.: Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs). RFC 4380, RFC Editor (February 2006), https://rfc-editor.org/rfc/rfc4380.txt
[30] IANA: IPv6 Global Unicast Address Assignments, https://www.iana.org/assignments/ipv6-unicast-address-assignments/ipv6-unicast-address-assignments.xhtml
[31] Internet2: Internet2, https://www.internet2.edu
[32] Karir, M., Huston, G., Michaelson, G., Bailey, M.: Understanding IPv6 Populations in the Wild. In: In Proceedings of the 14th International Conference on Passive and Active Measurement. pp. 256–259. PAM 2013 (2013)
[33] Kastenholz, F., Partridge, D.C.: Technical Criteria for Choosing IP The Next Generation (IPng). RFC 1726, RFC Editor (December 1994), https://rfc-editor.org/rfc/rfc1726.txt

[34] Krishnan, S., Davies, E.B., Savola, P.: IPv6 Transition/Co-existence Security Considerations. RFC 4942, RFC Editor (September 2007). https://doi.org/10.17487/RFC4942, https://rfc-editor.org/rfc/rfc4942.txt
[35] Lepinski, M., Kent, S.: An Infrastructure to Support Secure Internet Routing. RFC 6480, RFC Editor (February 2012), https://rfc-editor.org/rfc/rfc6480.txt
[36] Mawatari, M., Kawashim, M., Bryne, C.: 464XLAT: Combination of Stateful and Stateless Translation. RFC 6877, RFC Editor (April 2013), https://rfc-editor.org/rfc/rfc6877.txt
[37] Nakibly, G., Arov, M.: Routing Loop Attacks using IPv6 Tunnels. In: In Proceedings of the 3rd USENIX Workshop on Offensive Technologies. pp. 1–7. WOOT 2009 (2009)
[38] Nikkhah, M., Guerin, R.: Migrating the Internet to IPv6: An Exploration of the When and Why. IEEE/ACM Transactions on Networking **24**(4), 2291–2304 (August 2016)
[39] Patel, C., Savola, P.: Security Considerations for 6to4. RFC 3964, RFC Editor (December 2004), https://rfc-editor.org/rfc/rfc3964.txt
[40] Savola, P.: Observations of IPv6 traffic on a 6to4 relay. ACM SIGCOMM Computer Communication Review **35**(1), 23–28 (January 2005)
[41] Schinazi, D., Pauly, T.: Happy Eyeballs Version 2: Better Connectivity Using Concurrency. RFC 8305, RFC Editor (December 2017), https://rfc-editor.org/rfc/rfc8305.txt
[42] StatCounter: Desktop Windows Version Market Share Worldwide, http://gs.statcounter.com/windows-version-market-share/desktop/worldwide
[43] Thaler, D., Draves, R., Matsumoto, A., Chown, T.: Default Address Selection for Internet Protocol Version 6 (IPv6). RFC 6724, RFC Editor (September 2012), https://rfc-editor.org/rfc/rfc6724.txt
[44] Thurrott, P.: Windows 10 Version 1803 Surges to 90 Percent Usage Share (2018), https://www.thurrott.com/windows/windows-10/176435/windows-10-version-1803-surges-to-90-percent-usage-share
[45] Townsley, M., Troan, O.: Pv6 Rapid Deployment on IPv4 Infrastructures (6rd) – Protocol Specification. RFC 5969, RFC Editor (August 2010), https://rfc-editor.org/rfc/rfc5969.txt
[46] Ullrich, J., Krombholz, K., Hobel, H., Dabrowski, A., Weippl, E.: IPv6 Security: Attacks and Countermeasures in a Nutshell. In: In Proceedings of the 8th USENIX Workshop on Offensive Technologies. p. 5. WOOT 2014 (2014)
[47] Vixie, P., Schryver, V.: DNS Response Policy Zones, https://dnsrpz.info