2011 Ninth Annual Communication Networks and Services Research Conference

The Good, the Bad, the IPv6

Matthew Dunlop*† Stephen Groat*† Randy Marchany† Joseph Tront*
*Bradley Department of Electrical and Computer Engineering
†Virginia Tech Information Technology Security Office
Virginia Tech, Blacksburg, VA 24061, USA
Email: {dunlop,sgroat,marchany,jgtront}@vt.edu

***Abstract*—As more network-capable devices are developed, it becomes easier for us to remain connected to our friends, colleagues, and even our homes. Unfortunately, it becomes easier for unintended third parties to remain connected to us as well. The impact of this has been limited by our proximity to the third party (i.e. the same local area network) – until now. The Internet Protocol version 6 (IPv6) includes a method for devices to automatically configure their own addresses. This technique relieves some administrative burden, but provides an opportunity to monitor users from anywhere in the world. We explain why this problem exists. We also provide proof of host tracking and monitoring. We then illustrate examples of how both well-intended and malicious parties can exploit autoconfigured addresses. We submit that the benefits of using autoconfigured addresses do not outweigh the privacy implications. Finally, we present some alternatives to resolve this issue before IPv6 is deployed globally.**

***Index Terms*—Anonymity, IPv6 Addressing, Privacy**

I. Introduction

Numerous reports [1], [2] indicate that the transition to the Internet Protocol version 6 (IPv6) is inevitable. The main factor pushing this transition is the lack of free, unclaimed addresses in the Internet Protocol version 4 (IPv4). IPv6 solves the address space issue by allocating 128 bits to the address as opposed to the 32-bit address in IPv4. This means that IPv6 provides more than $3 \cdot 10^{38}$ addresses while IPv4 only provides approximately four billion. Another way to look at this is that IPv6 provides over $5 \cdot 10^{28}$ addresses for every one of the 6.8 billion people in the World [3].

This immense address space creates an additional management burden on administrators. To mitigate this burden, IPv6 designers created a method to allow individual hosts to create their own IPv6 addresses. This method is called StateLess Address AutoConfiguration (SLAAC). Using this technique, hosts create their own host portion of the IPv6 address. This host portion is referred to as the interface identifier (IID). To simplify processing and provide a unique address, the IID is often comprised of the network interface's Media Access Control (MAC) address [4]. Unfortunately, using a host's MAC address turns a link local address into a global address. Since the MAC address is static, this allows a third party, whether malicious or not, a means of geographically tracking users and monitoring their traffic through their network enabled devices.

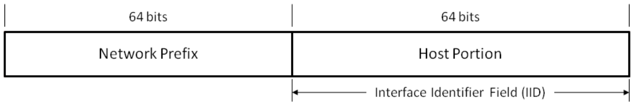

Fig. 1. IPv6 128-bit address format

There are a number of privacy ramifications that result from making MAC addresses part of a global static IID. Depending on the point of view, these ramifications can be spun as either good or bad. From a positive aspect, IID tracking brings a physical security aspect to cyber security. This convergence is referred to as converged security. Law enforcement officials can use static IIDs as a forensics aid. Unfortunately, there are negative aspects as well. Tracking people using IIDs is easily applied to cyberstalking. Additionally, terrorists can use IID tracking and monitoring to further their goals. Despite whether the intent is good or bad, over 70% of Americans in a 2006 survey were opposed to being monitored [5].

Our goal is to illustrate in detail how IPv6 SLAAC can be used for both altruistic and nefarious means. First, we provide a short background on IPv6 SLAAC in Section II. In Section III, we discuss similar work regarding tracking IPv6 addresses. We demonstrate that IPv6 address monitoring is possible in Section IV. A detailed explanation of some positive and negative implications of using static IIDs is outlined in Section V. We dedicate Section VI to explaining a few existing methods that help protect users' privacy. In Section VII we briefly cover some future work and conclude in section VIII.

II. Background

To handle the immense address space in IPv6, SLAAC was implemented. In this system, broadcast information provides nodes with half of their IPv6 addresses. The node uses operating system configuration parameters to create the other half, referred to as the IID (See Fig. 1). Some operating systems, such as Mac OS X and Linux, use the 64-bit Extended Unique Identifier (EUI-64) format to expand the address of the network adapter to fill the other half of the IPv6 address [4]. Other operating systems, such as Windows, deterministically obscure the network interface's address when the network card is installed [6]. Both these systems create a static IID. Assuring address privacy is necessary to protect systems from unwanted tracking and monitoring.

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Government, the Department of Defense, or any of its agencies.

978-0-7695-4393-2/11 $26.00 © 2011 IEEE
DOI 10.1109/CNSR.2011.20

A. Importance of Privacy

According to Westin, "Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others" [7]. The key point in Westin's definition is for individuals to determine for themselves what information is private. As society continues to migrate more information onto computers and other networked devices, information becomes more vulnerable to interested third parties. The result is that the decision of what information is private is being taken away from the individual. As new technologies are developed, the potential privacy implications need to be considered. IPv6 addressing is one example of a technology where convenience outweighed privacy.

B. Future of the Internet

The larger address space of IPv6 encourages new classes of devices to be networked, both mobile and stationary. A side effect of this trend is the creation of more data and attack vectors for malicious users to exploit. The proliferation of network devices in residential homes, such as appliances, produces more data about inhabitants' daily habits and potentially violates their privacy. For example, a networked dishwasher may be programmed only to run when the house is empty. If burglars can monitor global traffic from the dishwasher, they could determine the optimal time to burglarize a house. Also, by interconnecting mobile devices, such as cars, location information could be gathered by malicious users. For example, attackers could monitor vehicles for downloaded navigation maps. All sorts of disreputable plans could be made to greet the occupants once they arrive at that destination. The military has even discussed networking all munitions [8]. If the military begins using IPv6 IIDs to communicate with ordinance and other assets, the static nature of the IID could allow assets to be tracked and attacked after initial identification.

C. Privacy of Static IID

The static IIDs used in IPv6 addressing compromise a user's privacy. Creating a static IID from a MAC address through the EUI-64 expansion format [4] allows nodes to be logically and geographically tracked as they travel to different networks. Since the EUI-64 format results in a deterministic IID, users can be tracked on a network by scanning different subnets and searching for the MAC-generated IIDs. Using simple commands, such as ping and traceroute, the location of a user can be determined with reasonable geographic accuracy. Even the Windows obscuration of the IPv6 IID does not protect a user. Windows systems use a static IID for neighbor solicitation. Once this address is captured locally, it can be used globally to track a host using the same technique of searching subnets as used for unobscured host addresses.

By monitoring the traffic on a network over an extended period of time, a single user's traffic can be identified and analyzed. Armed with this data, a third party (whether malicious or not) can potentially tie a device to its actual user. As the user crosses different subnets, traffic can be collected and correlated by examining the static IID. This vulnerability to tracking does not typically apply when using IPv4. Most medium to large IPv4 networks implement the Dynamic Host Configuration Protocol (DHCP), which changes user addresses nondeterministically. As a result, DHCP logs are needed to tie traffic sniffed from a network with a particular user. Due to the deterministic IID in IPv6 address autoconfiguration, simple filters could be created to filter the traffic of a single user on any subnet. This would allow an interested party to identify and monitor a user's on-line activity through traffic analysis. In a dual-stack implementation where a node uses a mixture of IPv4 and IPv6, special Internet Control Message Protocol version 6 (ICMPv6) Neighbor Solicitation messages can even provide an interested party the IPv4 address linked with an IPv6 address. This correlation allows for traffic collection to extend to IPv4 for a single session.

III. Related Work

The idea that IPv6 SLAAC could lead to potential privacy concerns has been mentioned in a few Requests for Comment (RFCs). No one thus far has discussed how a third party could exploit a user's privacy though linking an IPv6 IID to a user. Further, no work has been done that discusses the implications that can result from this privacy violation. There has also been some work related to privacy issues concerning Mobile IPv6. These concerns are separate from those we focus on, but we briefly mention them for completeness.

Other researchers have identified that static IIDs in IPv6 can be used to track nodes. Narten et al. discussed this problem in RFC 4941 and concluded that a non-changing interface would allow an eavesdropper to correlate unrelated information with a particular node [6]. Haddad addressed the fact that mobile nodes using IPv6 SLAAC can reveal their location to an eavesdropper [9]. Our work builds on these ideas by discussing and demonstrating how an interested party can eavesdrop on a user from anywhere on the Internet using basic network tools. Additionally, we test this theory on a more than 30,000 node production IPv6 network. We then illustrate some specific applications, both good and bad, in an attempt to clarify just how powerful this technique can be.

There is also research that addresses issues related to potential privacy problems with regards to Mobile IPv6. Koodli discusses how mobile nodes' home or care-of addresses can be used to reveal that they have roamed [10]. Castelluccia et al. and Qiu et al. also discuss how mobile nodes can be tracked using home and care-of address [11], [12]. While in principle these concepts relate to privacy concerns with tracking of IPv6 node location, they focus on a completely unrelated vulnerability. Additionally, the vulnerability we address affects both mobile and stationary IPv6 nodes. Although Mobile IPv6 vulnerabilities are important, privacy concerns associated with the standard IPv6 must be addressed before Mobile IPv6 can be secured.

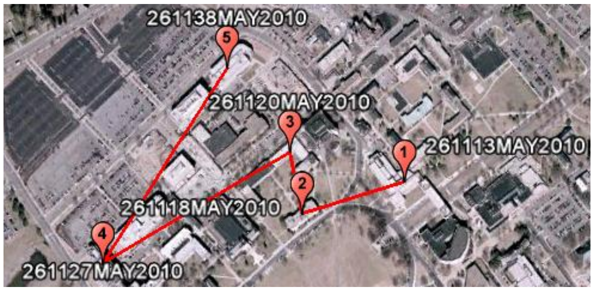

Fig. 2. Geotemporal plot of a wireless node's movement within the campus network

IV. Validation

We were able to validate that IPv6 address tracking and monitoring is possible using our campus-wide IPv6 production network. The network supports more than 30,000 IPv6 nodes on a daily basis. We conducted testing using an Android mobile device. The Android operating system uses the EUI-64 address format to form wireless IPv6 addresses.

The first part of our testing involved tracking the mobile device as it moved around campus. Geotemporal tracking was possible since the campus network is designed to have different subnets cover different geographic areas. We were constrained to campus due to the lack of IPv6 capable networks outside of campus. To conduct the test, we programmed a script that continuously sent echo requests to the different subnets on campus. When an echo reply was returned, we stored the time and location of the reply. Fig. 2 demonstrates the results of a successful tracking attempt.

The second part of our testing involved traffic monitoring. Our goal was to demonstrate that we could isolate a node, regardless of subnet, and collect all of its associated network traffic. A sensor was placed at the network border to collect all IPv6 traffic leaving the network. With the use of a packet sniffer, we were able to filter the traffic related to the node in question. Our filtered contents reflected all of the subnet locations illustrated in Fig. 2.

V. Implications

The ability to track users based off their stateless IPv6 addresses results in numerous consequences. Whether tracking is seen as helpful or malicious is a matter of perspective and application. We first present the positive aspects of IID tracking in Section V-A as a demonstration that not all applications are necessarily malicious. We then present what we consider to be negative applications of IID tracking in Section V-B. Regardless of the application, however, we assert that the positive aspects of IID tracking do not outweigh an individual's right to privacy.

A. The Good

There are many applications that network administrators and law enforcement officials would argue provide positive results of having the capability to track users. One such application is referred to as converged security. This term refers to the ability to combine physical security with cyber security. By tracking a user's physical location, administrators are afforded with an additional layer of security. Meanwhile, the numerous forensic possibilities provided by IID tracking would appear attractive to law enforcement officials.

1) Converged Security: As threats against organizations and individuals evolve in scope and complexity, physical and logical security systems must converge to create a holistic approach to security [13]. By leveraging all pieces of an enterprise's infrastructure in creating a security solution, security policies can be effectively enforced and valuable intelligence can be collected to ensure the safety of the workers and the enterprise. Individuals can also deploy the same technologies on a smaller scale to protect their families and homes. Converged security expands the resources available to a security system, allowing for logical systems to improve physical security.

IID analysis could help an organization enforce current physical and logical security policies through monitoring of network resources. Using IID analysis could supplement current control measures and provide an extra layer of security to protect valuable resources. Physically, access control measures could be supplemented by IID tracking. If an attacker followed a valid user through a physical control access point to "tailgate" off of their access, physical access control would rely on either the user or security guards to recognize and catch the tailgater. Using IID analysis, a wireless access control point could be created which checks the IID of a user's wireless device. If not in the database, the user would be flagged and security would be alerted. Logically, file access could be logged and traced through IID analysis, providing a system administrator an extra safeguard to protect information. For example, if sensitive data was only supposed to be accessed by specific systems, IID analysis could be used by a logical control system to assure that only authorized systems access resources. Incorporating IID analysis into logical control systems adds another layer of defense for administrators.

In emergency situations, detailed tracking of resources and IIDs allows for an individual's location and building's human density to be established, allowing for a targeted response. Determining who is in the building, where they are located, and where the population density is greatest in a disaster situation allows for first response teams to react with the proper resources to save the most lives possible. In a building fire, firefighters put themselves at risk searching for victims. Often, firefighters are exposed to unnecessary danger by searching unoccupied rooms. Crucial time is also wasted by this methodical search technique. By using IID analysis, the largest groups of victims could be identified to save the most possible lives. After the large groups have been evacuated, individuals could be quickly located through IID tracking. IID analysis can also be used to pinpoint those victims with special needs that might require special extraction equipment.

Similar to disaster scenarios, IID tracking could assist law enforcement during hostage situations. Law enforcement could use the IIDs of wireless devices within a building to determine

The Good

In a hostage situation, the physical locations and identities of hostages can be ascertained to determine how many hostages there are and physically where they are located.

A forensics team could geotemporally track the location of a stolen device or a suspected criminal using globally accessible IIDs traced from a crime scene.

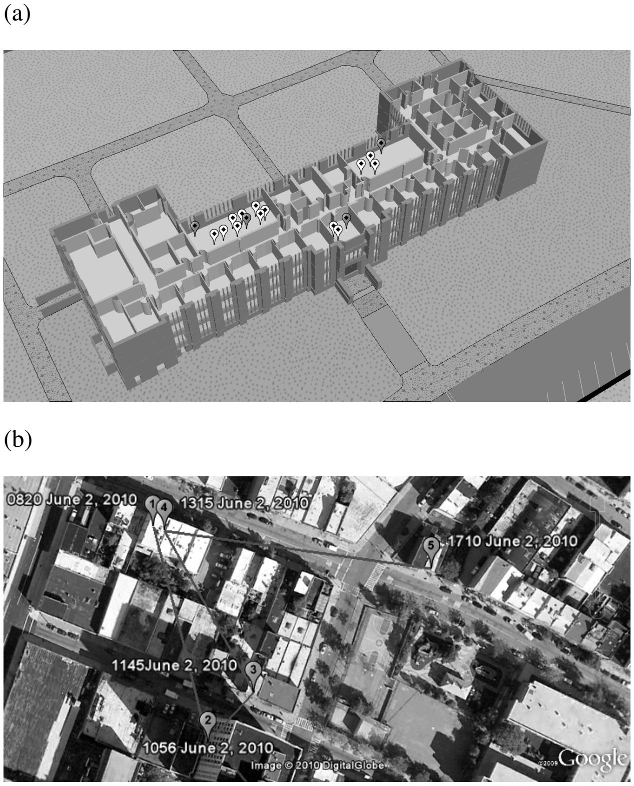

The Bad

Terrorists could track the number of distinct IIDs within a target structure before initiating an attack.

A cyberstalker can track a target's daily activities and predict a movement pattern.

Fig. 3. Sample scenarios where IPv6 IID tracking can apply. Captions on the left describe positive applications of each depicted scenario, while captions on the right describe negative applications.

the number, identities and approximate locations of hostages. The approximate number of captors and their locations might even be determined by looking at unknown IIDs within the structure. Armed with this information, SWAT teams can develop rescue plans that will minimize casualties. An example of this type of scenario is depicted in Fig. 3(a), where white pins indicate the locations of hostages and gray pins indicate captor locations.

Personal and family security could also be enhanced through IPv6 tracking capabilities. Currently, many cellular phone service providers offer a locater services, such as a Global Positioning System (GPS), in phones to report the location of a unit. Since most users always carry their cell phones and have them powered up, this is a viable way of tracking people. While simple to implement and deploy, using a cellular phone as the only tracking device creates a single point of failure. If the phone is not in tower range, runs out of battery, or is turned off, the user cannot be tracked. As more devices become interconnected, the number of IPv6 addresses and IIDs associated with an individual will also increase. By monitoring all of the devices that have static or deterministic IIDs, tracking is moved from a single device to all of the Internet-capable devices that a user deploys daily. Since IID tracking is not reliant on GPS and can be used on any Internet-capable device, more of the devices that people carry can serve as tracking sources.

2) Forensics: While some disciplines are faced with excessive amounts of information, forensic investigations often suffer from a lack of data when trying to solve a crime or identify a criminal. Using IID analysis to identify or exonerate possible suspects provides investigators with another information source. Gleaning publicly available information, such as geotemporal location from IID analysis, may help investigators determine a suspect's location at the time of a crime. Also, geotemporal IID tracking could allow for the tracking and the recovery of stolen networked assets, such as cars or laptops as illustrated in Fig. 3(b). Finally, when trying to reconstruct the scene of a crime, IID analysis could tie a suspect to the crime scene.

Using IID analysis to determine a person of interest's location at the time of a crime could provide law enforcement with additional evidence to convict a suspect. Currently, many criminal cases rely on witness testimony to place an individual at a crime scene. Witness testimony is often inaccurate and unreliable and, therefore, discredited. Yet, scientifically supported evidence, such as fingerprints or DNA, often convinces juries and proves innocence or guilt. By logging wireless network access information, a wireless device could be used to implicate its user. Supplementing testimony with an IID gleaned from saved data provides scientifically supportable data for evidence.

To track and recover stolen property, including cars and laptops, IID analysis could be used to geographically track networked assets. While asset tracking systems are currently deployed, the proprietary systems require specialized hardware and are dependent on network implementation. For example, the current LoJack® system used to track stolen vehicles requires a specialized transmitter and uses a proprietary radio system. Additionally, it could be subject to limited range in certain environments. As cars become networked, tracking systems integrated into the car's design could be embedded is such as way as to increase reliability and make removal difficult. This would prevent attackers from removing the transmitter and expand the coverage and reliability of the network.

Forensic examination of access logs and network traffic could also be supplemented by IID analysis. When logs collect information about an address that accessed the system, static IIDs provide information about the connecting system. Failed logon attempts could be correlated with specific users, discouraging unauthorized access attempts. Monitoring of network traffic by the government may also be expanded in the near future. In 2003, a petition for expedited rule making was made to the Federal Communication Commission (FCC) on behalf of the Department of Justice (DoJ) and its investigatory agencies to allow all wireless header information to be sniffed without warrants [14]. Assuming the use of IPv6 SLAAC, legally sniffing all header information would allow for the DoJ to track any user's location and traffic. This information would facilitate the cost-effective and manpower-efficient tracking of people of interest.

B. The Bad

It is not difficult to imagine that the capability to track users can be exploited by malicious nodes for nefarious means. One of the more obvious malicious applications is cyberstalking. On a larger scale, terrorists can take advantage of the ability to track people and assets to further their goals without alerting authorities to their intentions.

1) Cyberstalking: There are many different definitions for cyberstalking. Most definitions differ by the means used to stalk a victim. Some stalkers use email while others eavesdrop on their victim's communications. Regardless of how the stalking is done, cyberstalking refers to the repeated unwanted attention of a stalker using some sort of electronic means [15].

Cyberstalking executed through common methods involves an active effort on the part of the stalker. Additionally, there is a level of expertise required. The average cyberstalker is not able to hack into a victim's machine or even place and run a packet sniffer. Assuming that the cyberstalker possesses the skills to accomplish these tasks, there is the risk of the intrusion being traced back to the cyberstalker. There is also the possibility that the victim has a secured machine, making penetration difficult. A mobile victim makes the cyberstalker's task more difficult. As a victim moves between subnets, the DHCP address provided is not logically connected to the victim. If the cyberstalker does not have a physical presence on the victim's machine, tracking the victim becomes much more difficult.

SLAAC in IPv6 alleviates many of the challenges and risks for a cyberstalker. Since the IID is static, the cyberstalker always knows half of a victim's Internet Protocol (IP) address. The other half, the subnet portion, is easily discovered by conducting a thorough reconnaissance of the geographic areas frequented by the victim. Armed with the IID and subnets, the cyberstalker can continually ping the likely subnets for the victim's IID. A successful ping reply indicates that the victim is at that specific location. This form of attack will not alert the victim, yet will provide the cyberstalker with the victim's daily movements. The attacker may even be able to establish a movement pattern for the victim as illustrated in Fig. 3(b). This pattern could be used to plan a physical assault, burglary, or other crime. The proliferation of handheld network-capable devices aids a cyberstalker's ability to keep constant tabs on victims.

Monitoring a victim's network traffic is also made easier by IPv6 SLAAC. A cyberstalker running a packet sniffer on a congested network does not need to perform tedious traffic analysis to filter out the victim's traffic. Since the IPv6 IID is static, the cyberstalker can easily set up a filter rule to exclude all traffic not matching the victim's IID. This provides the cyberstalker with only traffic belonging to the victim. DHCP in IPv4, on the other hand, issues IP addresses to victims that change based on lease duration. These DHCP addresses are unassociated with the hosts' identities. Traffic monitoring in IPv4 is not nearly as straightforward as it is when using IPv6 autoconfigured addresses.

Tracking victims or monitoring their network traffic may not even be the goal of the cyberstalker. It may be that the cyberstalker wants to cause pain and humiliation to the victim. SLAAC in IPv6 provides the cyberstalker with this capability. Since the IID is static and linked to a specific device, a cyberstalker can pretend to be the victim. This form of attack works when the victim is being monitored by a third-party, such as an employer. By spoofing the victim's IID, a cyberstalker can visit locations of ill-repute or conduct illegal Internet activities. Although the victim can likely prove his/her innocence through computer logs, the initial accusation followed by the subsequent discomfort felt by the victim accomplishes the cyberstalker's goals.

2) Terrorism: A potentially disastrous application of IPv6 SLAAC is toward forwarding the goals of terrorists. The Code of Federal Regulations defines terrorism as "the unlawful use of force and violence against persons or property to intimidate or coerce a government, the civilian population, or any segment thereof, in furtherance of political or social objectives" [16]. The best way to cause fear is through surprise and shock value. Terrorists can take advantage of the tracking capabilities provided by static IIDs to plan and organize attacks without attracting attention.

Since users can be easily identified through their IIDs, terrorists can use the IIDs to target individuals. A terrorist can use traffic analysis as described in Section V-B1 to pair a target with his/her IPv6 IID. The terrorist can then use cyberstalking techniques to develop a movement pattern for the target. Once a pattern is discerned, the terrorist can choose the most suitable location to trap the target and wait. If the terrorist's intention is assassination, the chosen location can be set up to remotely detonate an explosive device when the victim's IID is detected. If the goal is fear, a similar trap can be sprung remotely. In either case, the terrorist can be miles away and evade capture. This type of attack would be difficult to detect and has the potential to cause widespread panic.

A terrorist can also use IID tracking in an attempt to undermine an organization. By monitoring the locations and network traffic of individuals within an organization, a terrorist can apply social network analysis principles to determine the key people within an organization. There are certain people, not necessarily the leadership, within an organization that are critical to the functionality of the organization [17]. Once these people are identified, a terrorist can target each person individually. If these people are removed, the organization will falter and possibly fail.

Another way terrorists can take advantage of IID tracking is in the planning of attacks against soft targets. Soft targets are typically nonmilitary targets that are unarmored and/or undefended. Terrorist attacks are usually planned against soft targets since terrorists are often ill-equipped to confront military targets. Additionally, explosives targeted against soft targets generate higher casualties and, therefore, more media attention [18]. Attacks against soft targets usually require a great deal of planning. During the planning phase, terrorists can use packet sniffers to collect all the distinct IIDs that access the network of the target. If the number of victims is the terrorists' goal, then who owns the IIDs is unimportant. Armed with a list of IIDs, the terrorists can run a script to ping each IID and maintain a counter. Once the counter reaches a predetermined threshold (see Fig. 3(a)), the attack can be launched. The terrorists need not be anywhere within the vicinity of the target. This type of tracking is not possible in IPv4 due primarily to the use of network address translation (NAT), which hides the number of nodes on a subnet. DHCP in IPv4 also makes this type of attack difficult. Since IPv4 addresses are issued nondeterministically, a terrorist would have to physically scan the entire subnet to determine the number of active hosts. Depending on the size of the network, this could take time.

Terrorists can also track locations of law enforcement to plan the initiation of an attack or hostage situation. Law enforcement vehicles will each have unique identifiable IPv6 IIDs. Terrorists could take advantage of this by collecting the IIDs tied to all law enforcement vehicles operating within a particular region. Since law enforcement vehicles communicate wirelessly, collecting these IIDs is trivial if a thorough reconnaissance is conducted. Armed with the IIDs of area law enforcement, terrorists can initiate an attack when there are no law enforcement in the proximity of the attack site. Additionally, terrorists can anticipate the response time required for law enforcement to arrive.

Terrorists can take advantage of IPv6 IIDs to track military troop movements and size. Since military units are becoming digitized, units and possibly even individual soldiers may communicate over IPv6 wireless networks. Terrorists could use the IIDs to identify specific units. Through traffic analysis they could identify which units are communicating with one another. They could possibly even predict future operations. If individual soldiers are communicating over an IPv6 network, terrorists can pinpoint and track key leaders as well as unit strength.

Using a device's MAC address to compose the IID is another vulnerability that terrorists can exploit. The MAC address is composed of two parts. The first half is called the Organizational Unique Identifier (OUI) and identifies the vendor that issues the Network Interface Controller (NIC). The second half is assigned by the vendor and is designed to act like a serial number to make the MAC unique [19]. Usually, agencies will contract a vendor to produce all of a specific product. This means that, with high probability, an entire product line shares the same OUI. Terrorists could use this to locate all of a specific type of asset. Minimally, a terrorist could learn how many assets there are and possibly their locations. More importantly, a terrorist could identify a vulnerability specific to a particular brand of asset. The terrorist could then use the OUI to limit the search for those assets within an organization and launch an exploit against the vulnerability. For example, the Smart Grid is used to deliver electricity to consumers [20]. If Smart Grid devices were vulnerable to a particular exploit, terrorists could target the exploit against all systems sharing the same OUI. Terrorists in control of the Smart Grid could cause a host of problems, such as shutting down electrical power to a geographic area.

VI. Privacy Protection

Regardless of the intent behind IID tracking, users should be entitled to an expectation that their privacy will be maintained. The argument exists that if a person has nothing to hide, then they should not care if they are monitored [21]. Aside from the malicious possibilities we mentioned in Section V-B, most people just prefer privacy. When a homeowner constructs a privacy fence in their backyard, are we to assume that it is because they want to participate in some sort or illicit or illegal behavior? Similarly, if we do not plan to discuss illegal activities, should we agree to have all our phone conversations recorded? The answer to both of these hypothetical scenarios is emphatically, “No.” Therefore, it is important to prevent IID tracking before IPv6 is globally deployed.

To that end, there are three different approaches to preventing IID tracking in IPv6. The first approach involves obscuring the IID. There are two proposals for IID obscuration. The first obscuration approach uses cryptographically generated addresses [22] while the second obscuration approach uses what are called privacy extensions [6]. Another way of preventing IID tracking is through the use of the Dynamic Host Configuration Protocol for IPv6 (DHCPv6) [23]. The final method is through the use of the Encapsulating Security Payload (ESP) for Internet Protocol Security (IPsec) [24].

The first of the two IID obscuration techniques uses Cryptographically Generated Addresses (CGAs). In general, a CGA is formed by hashing the sender's public key along with some other parameters [22]. The original purpose of CGAs was to prevent denial of service attacks against the SEcure Neighbor Discovery (SEND) protocol [25]. Since CGAs also dynamically obscure IPv6 autoconfigured addresses, they can be applied as a defense against IPv6 address tracking.

The main disadvantage to using CGAs is the computational cost. Producing an acceptable CGA involves the generation of two hash values, Hash2 followed by Hash1. The complexity of generating Hash2 depends on the strength of a security parameter (*Sec*). The security parameter can take on any value from 0-7 and indicates the number of leading zeros Hash2 must contain. The number of zeros is determined by multiplying *Sec* by 16. On average, it takes O($2^{16 \cdot Sec}$) iterations to generate Hash2. Once an acceptable Hash2 is computed, Hash1 is generated using some of the final Hash2 parameters as well as the subnet prefix. The leftmost 64 bits of Hash1 are used as the IID with the exception of five bits used for other purposes [22]. At this point, duplicate address detection is conducted [25]. If three duplicate addresses are detected, the IID is rejected and the process starts anew. The large number of hash calculations required to generate a CGA could quickly overwhelm a power-constrained device.

Privacy Extensions provide the second of the two IID obscuration techniques. Privacy extensions generate a random IID by hashing the concatenation of a user's EUI-64 IID with a 64-bit "history value" and taking the leftmost 64 bits. The "history value" is initially produced from the leftmost 64 bits of a pseudo-random number. From this point, "history values" are produced using previously calculated IIDs. Using "history values" instead of pseudo-random numbers for each IID calculation limits the number of duplicate address collisions that occur due to only using 64 bits of the hash. If a duplicate address is detected, a new "history value" is formed and the process is repeated [6].

The disadvantages of using privacy extensions are not as severe as those of using CGAs. Assuming no address collisions, only one hash calculation is required of the sender to produce an obscured IID. Privacy extensions also carry parameters to limit the time an obscured IID remains valid. Unfortunately, the default values of these parameters are set too long. It is feasible for an IID using privacy extensions to remain static for as long as one week. During this time period, a malicious node could still successfully profile a target host. Fortunately, RFC 4941 allows users to modify these defaults [6].

The main factor that makes IID obscuration an attractive solution for hiding IPv6 addresses is that there is no need for any management overhead. Obscuration and verification both occur at the respective end hosts with no intervention by a trusted third party required. Although CGAs use a public key, the key is self-generated by the sender [22]. Privacy extensions use a history value that is generated based on a pseudo-random number. This lack of management makes IID obscuration scalable.

The Dynamic Host Configuration Protocol for IPv6 [23] is not an obscuration technique but rather a means to provide stateful address configuration. In IPv6, DHCPv6 works very similar to DHCP in IPv4. Instead of clients generating their own addresses, a DHCP server leases temporary addresses to clients. The main advantage of this technique is that the addresses issued by the DHCP server are typically not tied to the identity of a client. In principal, each time a client connects to a network, the DHCP server could issue a new address. Unfortunately, this does not happen in practice. RFC 3315 promotes the issuance of non-temporary addresses to clients. Clients do have the ability to request temporary addresses, which mask their location and activities globally. Locally, however, an attacker can still track clients through a DHCP Unique Identifier (DUID) that gets communicated between the client and server. The scope of this method of tracking is limited to the subnet of the client, server, or any relays [23]. There is also an administrative management burden that accompanies the use of DHCPv6.

IPsec also provides a means to protect users from tracking. In IPsec, this is accomplished through the use of the ESP in tunnel mode. The reason this hides the identity of the target node from being tracked is that the target node's entire packet, including address, gets encrypted. This encrypted portion then becomes part of the payload of a new packet using the address of the tunnel start point instead [24]. Of course, the tunnel start point cannot be the same as the target host or tracking will again be possible. One major advantage to using IPsec in tunnel mode is that the cryptographic burden of encryption and decryption is offloaded to the tunnel endpoints. This is extremely beneficial for power constrained devices.

There are, however, severe disadvantages to using IPsec as a privacy protection mechanism. The most striking is that IPsec used in this way requires a global key management infrastructure that does not currently exist [26]. Another disadvantage is that IPsec in tunnel mode only protects target nodes from those nodes external to the tunnel. Nodes residing on the same subnet as either tunnel endpoint will still be able to track the target nodes. This may provide a slight obstacle to the majority of malicious nodes, but will provide no obstacle to administrators. Depending on the point of view, this could be seen as either positive or negative.

While these prevention techniques seem simple and easy to implement, there are obstacles to implementation. The extra overhead and decreased performance required for frequent hash calculations [27] in embedded devices has caused IID obscuration to be ignored. DHCPv6 will not be implemented any time soon due to the additional network configuration and equipment needed. Also, since the lack of DHCP is touted as a feature of IPv6, few network administrators are choosing to implemented the service. IPsec suffers obstacles to global deployment because of the requirement for a global key management infrastructure. As more embedded devices become Internet capable, a user's identity becomes easier to determine. This is especially true as more attributes of a user are sent over the Internet. Regardless of the obstacles, some method of prevention should be implemented now. Of those mentioned, obscuration through the use of privacy extensions seems to be the best option in terms of benefits versus computational cost.

VII. Future Work

The next phase of our research will focus on DHCPv6. We plan to demonstrate that the DUID can be used to identify and monitor users. The DUID does not exist in DHCP for IPv4 and is used by the IPv6 protocol to identify a client to a server and vice versa [23]. Although the DUID does not have the same global scope as IIDs used in SLAAC, the scope is large enough to be of concern. In order to identify users through their DUIDs, the malicious node can be on the same subnet as the client, the DHCP server, or any of the DHCP relays. In the case of DHCP, we classify the person doing the tracking as being malicious because an administrator would already have the host identity simply because the host is issued the address by the administrator.

We also plan to design and implement an address obscuration technique. Each of the techniques described in Section VI has associated shortcomings. The main shortcoming is computational complexity. Our technique will be designed with the goal of minimizing computational complexity, thus making IPv6 address obscuration feasible for power-constrained devices. We also plan to test the validity and overhead of our design using our campus-wide IPv6 production network.

VIII. Conclusion

The Internet Protocol version 6 provides many benefits over version 4. Foremost among those is the much needed increase in address space. However, mitigating the administrative burden of managing this immense address space had lead to some oversights. SLAAC, in its current state, makes it too easy for a third-party to monitor the activities of unsuspecting targets. There are some powerfully compelling positive applications that IID tracking could help with. Unfortunately, there are also devastatingly negative applications. Fortunately, IPv6 is not yet fielded globally. Now is the time to design a solution which prevents the ability to track individuals through their IPv6-enabled devices. If nothing is done, we will all feel the effects of *the good, the bad, the IPv6*.

References

[1] J. Leyden, "Net shakeup looms as IPv4 resources start running low," The Register, Jun. 2010. [Online]. Available: http://www.theregister.co.uk/2010/06/01/ipv4_exhaustion_analysis/

[2] I. Lazar, "Is IPv6 in your future?" May 2010. [Online]. Available: http://www.networkworld.com/community/node/61346

[3] "U.S. & World population clocks," Available at http://www.census.gov/population/www/popclockus.html/ accessed 4 Mar. 2010.

[4] R. Hinden and S. Deering, "IP Version 6 Addressing Architecture," RFC 4291 (Draft Standard), Internet Engineering Task Force, Feb. 2006. [Online]. Available: http://www.ietf.org/rfc/rfc4291.txt

[5] L. Ponemon, "Americans' perceptions about surveillance," Mar. 2006.

[6] T. Narten, R. Draves, and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6," RFC 4941 (Draft Standard), Internet Engineering Task Force, Sep. 2007. [Online]. Available: http://www.ietf.org/rfc/rfc4941.txt

[7] A. Westin, *Privacy and Freedom*. New Jork Atheneum, 1967.

[8] C. E. Croom, "C4I integration - operational advantage for the warfighter," Sep. 2003.

[9] W. Haddad, "Privacy for mobile and multi-homed nodes: MoMiPriv problem statement," United States, 2005.

[10] R. Koodli, "IP Address Location Privacy and Mobile IPv6: Problem Statement," RFC 4882 (Informational), Internet Engineering Task Force, May 2007. [Online]. Available: http://www.ietf.org/rfc/rfc4882.txt

[11] C. Castelluccia, F. Dupont, and G. Montenegro, "A simple privacy extension for mobile IPV6," in *Mobile and Wireless Communication Networks, IFIP TC6 / WG6.8 Conference on Mobile and Wireless Communication Networks (MWCN 2004)*, Oct. 2004, pp. 239–249.

[12] Y. Qiu, J. Zhou, F. Bao, and R. Deng, "Protocol for hiding movement of mobile nodes in Mobile IPv6," in *62nd IEEE Vehicular Technology Conference*, vol. 2, Sep. 2005, pp. 812 – 815.

[13] B. T. Contos, C. Derodeff, W. P. Crowell, and D. Dunkel, *Physical and Logical Security Convergence: Powered By Enterprise Security Management*. Syngress, 2007.

[14] Federal Bureau of Investigation, "Electronic surveillance needs for carrier grade voice over packet (CGVoP) service," Jan. 2003, cALEA Implementation.

[15] Attorney General to the Vice President, "Cyberstalking: A new challenge for law enforcement and industry," 2010.

[16] National Archives and Records Administration, *Code of Federal Regulations, Title 28, Section 0.85*. U.S. Government Printing Office, Jul. 2001, vol. 1.

[17] I. McCulloh, "Detecting changes in a dynamic social network," Ph.D. dissertation, Carnegie Mellon University, Mar. 2009, cMU-ISR-09-104.

[18] STRATFOR Global Intelligence, "The terrorist attack cycle: Selecting the target," Sep. 2005.

[19] R. Deal, *Cisco Certified Network Associate Study Guide (Exam 640-802)*. The McGraw-Hill Companies, 2008.

[20] U.S. Department of Energy, "Smart grid," Available at: http://www.oe.energy.gov/smartgrid.htm accessed 9 Jun. 2010.

[21] D. J. Solove, "'I've got nothing to hide' and other misunderstandings of privacy," *Social Science Research Network Working Paper Series*, Jul. 2007.

[22] T. Aura, "Cryptographically Generated Addresses (CGA)," RFC 3972 (Proposed Standard), Internet Engineering Task Force, Mar. 2005, updated by RFCs 4581, 4982. [Online]. Available: http://www.ietf.org/rfc/rfc3972.txt

[23] R. Droms, J. Bound, B. Volz, T. Lemon, C. Perkins, and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," RFC 3315 (Proposed Standard), Internet Engineering Task Force, Jul. 2003, updated by RFCs 4361, 5494. [Online]. Available: http://www.ietf.org/rfc/rfc3315.txt

[24] S. Kent and R. Atkinson, "IP Encapsulating Security Payload (ESP)," RFC 2406 (Proposed Standard), Internet Engineering Task Force, Nov. 1998, obsoleted by RFCs 4303, 4305. [Online]. Available: http://www.ietf.org/rfc/rfc2406.txt

[25] J. Arkko, J. Kempf, B. Zill, and P. Nikander, "SEcure Neighbor Discovery (SEND)," RFC 3971 (Proposed Standard), Internet Engineering Task Force, Mar. 2005. [Online]. Available: http://www.ietf.org/rfc/rfc3971.txt

[26] A. Choudhary, "In-depth analysis of IPv6 security posture," in *The 5th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom 2009)*, Nov. 2009, pp. 1 –7.

[27] M. Durvy, J. Abeillé, P. Wetterwald, C. O'Flynn, B. Leverett, E. Gnoske, M. Vidales, G. Mulligan, N. Tsiftes, N. Finne, and A. Dunkels, "Making sensor networks IPv6 ready," in *SenSys '08: Proceedings of the 6th ACM conference on Embedded network sensor systems*, New York, NY, USA, 2008, pp. 421–422.