GCA Internet Integrity Papers:

Expanding IoT Honeypots to Include IPv6-Connected Devices

JULY 2024

GLOBAL CYBER ALLIANCE™

Foreword

I began working on IPv6 in the 1990s, and even then we knew that the number of Internet-connected devices was about to explode, that IPv4 addressing would be insufficient, and that transition mechanisms would require additional effort. IPv6 seemed an obvious choice in future development, but it never rose to the level of "committed" in development plans. It was often just below the line.

By 2008, IPv6 had seen major progress, but operational deployment was still lacking. At the Internet Society, I worked to facilitate IPv6 deployment in operational networks. Collaborative discussions led to "**World IPv6 Launch**."[1] Network operators made tremendous strides to get their services IPv6 capable in time, and I was happy to do my part in helping.

More than ten years later, Google reports that almost half the traffic they see on the Internet runs over IPv6. But the presence of IPv6 capabilities presents new opportunities for malicious actors to exploit vulnerabilities, and new challenges for network operators and owners of Internet-connected devices to protect themselves from malicious actors.

Now at the Global Cyber Alliance, I am working on ways to mitigate unwanted traffic, including operating a honeyfarm to detect the distribution of malware. This paper is where these worlds collide and we ask: how can we continue to expand the Internet with the endless possibilities of IPv6, and also protect users from cyber risks like malware?

July 9, 2024

Phil Roberts
Director of Technology
Global Cyber Alliance

[1] https://www.worldipv6launch.org/

GLOBAL CYBER ALLIANCE

Executive Summary

In this paper, we are specifically interested in stopping malware at its source by extending **ProxyPot**, the Global Cyber Alliance (GCA)'s proprietary honeypot technology, to detect attacks over SSH, Telnet, HTTP, and HTTPS, over both IPv4 and IPv6— a first of its kind.

It is relatively easy to scan the entire IPv4 Internet for threats, but IPv6 is on the rise, and it is so big we cannot use the same techniques we have been using on IPv4. IPv6-based attacks are already happening, and we must find ways to identify and mitigate them.

We explain GCA's AIDE project, recap the difficulty of scanning IPv6 address space for attacks, outline current IPv6 practices that might increase vulnerability, define best practices at the device level, and explore potential IPv6 scanning options.

We encourage you to check your own IPv6 resources and security practices and get involved with AIDE to support this work.

KEY TAKEAWAYS

- IPv4 remains pervasive, making it a much more common attack vector than IPv6– for now.
- With the increasing prevalence of IPv6, we can only expect IPv6-based attacks to grow.
- Attack detection is extremely limited over IPv6; we must detect and document attacks to stay ahead of the curve.
- We may be able to enhance ProxyPot to detect IPv6 attacks - a first of its kind.

Table of Contents

Honeypots are isolated networks that mimic valuable systems. They draw attackers so that the honeypot creator can monitor interactions and use that information to refine intrusion detection systems, improve threat responses, and better manage and prevent attacks.

Traditionally, honeypots have only been available using IPv4[2] as the underlying protocol. With the increasing prevalence of IPv6, we can only expect IPv6-based attacks to grow.

This paper investigates the possibility of including the much larger IPv6 address space in honeypot and honeyfarm operations.

Global Cyber Alliance (GCA)'s AIDE[3] project uses a global network of honeypots to detect, record, and analyze attacks against devices that look like Internet of Things (IoT) devices. Our honeyfarm is based on open-source software that detects attacks mainly through the SSH and Telnet protocols, and entirely using IPv4. Our proprietary software, ProxyPot, expands the detected attacks to include HTTP and HTTPS as attack channels.

It may be possible for this software to detect attacks that use IPv6, thus offering a honeypot that detects attacks over SSH, Telnet, HTTP, and HTTPS, over both IPv4 and IPv6—a first of its kind.

A honeypot is only useful if it is easily found by malicious actors. With today's computing power, it is relatively easy for an attacker to scan the entire IPv4 address space for vulnerable hosts, often checking against multiple ports at the same time, to identify an open port that can be explored for further attack.

[2] IPv4 was the first widely used protocol on the Internet, with just over four billion IP addresses. An updated IPv6 protocol includes over 340 undecillion IP addresses. Both IPv4 and IPv6 are in use on the Internet today.

[3] https://www.globalcyberalliance.org/aide/

This is not feasible in IPv6. First, its unimaginable size makes it impractical to fully scan, and second, finding one active host does not necessarily mean that there is another active host anywhere "nearby."

We investigate current practices in IPv6 operations that may make more limited scanning fruitful, and report on existing policy recommendations that may make IPv6 devices harder to find. We survey previous work by researchers interested in understanding the deployment landscape of IPv6, as those provide additional channels for understanding how IPv6 hosts may be found by third parties, including both researchers and malicious actors.

The cross-section of these findings leads us to possible approaches to deploy IPv6 honeypots in a way that lets malicious actors find them, and allows us to record attacks from those malicious actors. We conclude with best practices for address assignments and general security practices, and outline areas for further study.

Background

Even if the honeypots in GCA's AIDE network resemble IoT devices, over time, we have observed that many of the attacks collected are against generic Linux platforms, and some of the attacks actually look for higher powered devices rather than IoT devices before proceeding with downloading malware. Therefore, although the system is designed to detect attacks against IoT devices, it detects attacks against a wider set of platforms.

As stated in the **Introduction**, many attacks start by scanning the entire IPv4 address space for vulnerable ports and hosts. Such scans are now trivially simple and reasonably quick to execute, giving attackers the freshest available data about where to put effort into attacking. IPv6's vastness makes this nearly impossible.

IPv6 usage has grown substantially in the last decade. This is possible in part because host software providers have implemented IPv6 and enabled it by default on a range of devices. All major laptop and phone vendors, for instance, have done this implementation and so, when a network becomes IPv6 enabled, IPv6 connectivity occurs efficiently and transparently to the device owner. In addition, many devices will create and establish IPv6 connectivity between devices inside a home network even if IPv6 is not available by the external service provider.

While these devices provide both IPv4 and IPv6, the devices will remain vulnerable to attacks on their IPv4 interfaces. As long as IPv4 remains in operation and devices are easy to find, attackers will continue to focus on IPv4-based attacks over much harder IPv6-based attacks.

Having said that, we already observe hundreds of attacks against IPv6 devices every day. It is a small percentage of attacks compared to the thousands or tens of thousands against IPv4 interfaces, but these numbers are bound to grow as IPv6 becomes more available and IPv4 becomes more limited.

We must develop tools to detect and document the nature of IPv6 attacks to stay ahead of the curve— our efforts may be the first of their kind.

It is possible to envision devices becoming available —and perhaps growing in prevalence— that have only IPv6 connectivity. The Matter[4] initiative proposes to provide all device connectivity using IPv6 only. Gateways or gateway functionality will be provided in the near term to enable these IPv6-only devices to communicate from inside networks that have only IPv4 connectivity externally. As long as the gateway devices are not compromised, IPv6-only devices in IPv4-only networks will not be detectable.

Our ultimate interest is to determine how to protect IPv6-enabled devices from malicious activity. We would like to do this in part by using a network of IPv6-enabled honeypots to observe, record, and analyze IPv6 attacks. But, whereas it is easy to deploy an IPv4 honeypot and have it discovered by malicious actors, the above realities of the IPv6 Internet make it harder. So the immediate goal is to determine how to deploy an IPv6 honeypot, and then an IPv6 honeynet, in such a way that malicious actors will find and attack these devices.

Security and measurement researchers have already begun investigating how to find and identify devices over IPv6 connectivity. Much of this research is benign (how do we measure the IPv6 Internet?) or even intended to be beneficial (how do we provide security on the IPv6 Internet?). This is why part of the objectives of this research is to provide some early analysis of possible threat vectors a malicious actor might explore.

In this study, we begin with a brief survey of what has been reported and suggest some possible areas of further research based on our findings. We discuss how the practices of network operators, home router vendors, and endpoint software providers might contribute to some of the vulnerabilities and make suggestions about alternatives to reduce the threat surface.

Additionally, we suggest areas of further investigation and describe how GCA's AIDE project might be used to help inform best practices for policymakers to improve overall Internet security as the Internet evolves to more IPv6 deployment and even begins to experience IPv6-only deployments.

[4] The Matter initiative is operated under the aegis of CSA-IOT: https://csa-iot.org/all-solutions/matter/.

Terminology and Definitions

An **IPv4 address** is 32 bits and an **IPv6 address** is 128 bits. **IPv6 addresses** have two parts— a **network (subnet) prefix** of length n bits, and an **interface identifier (iid)** part of length 128-n bits. Regardless of the length of the subnet prefix, it is common to think of the interface identifier part of the address as 64 bits in length. The Internet Engineering Task Force (IETF)[5] Request for Comments (RFC) **RFC 4291**[6] describes IPv6 addresses in detail.

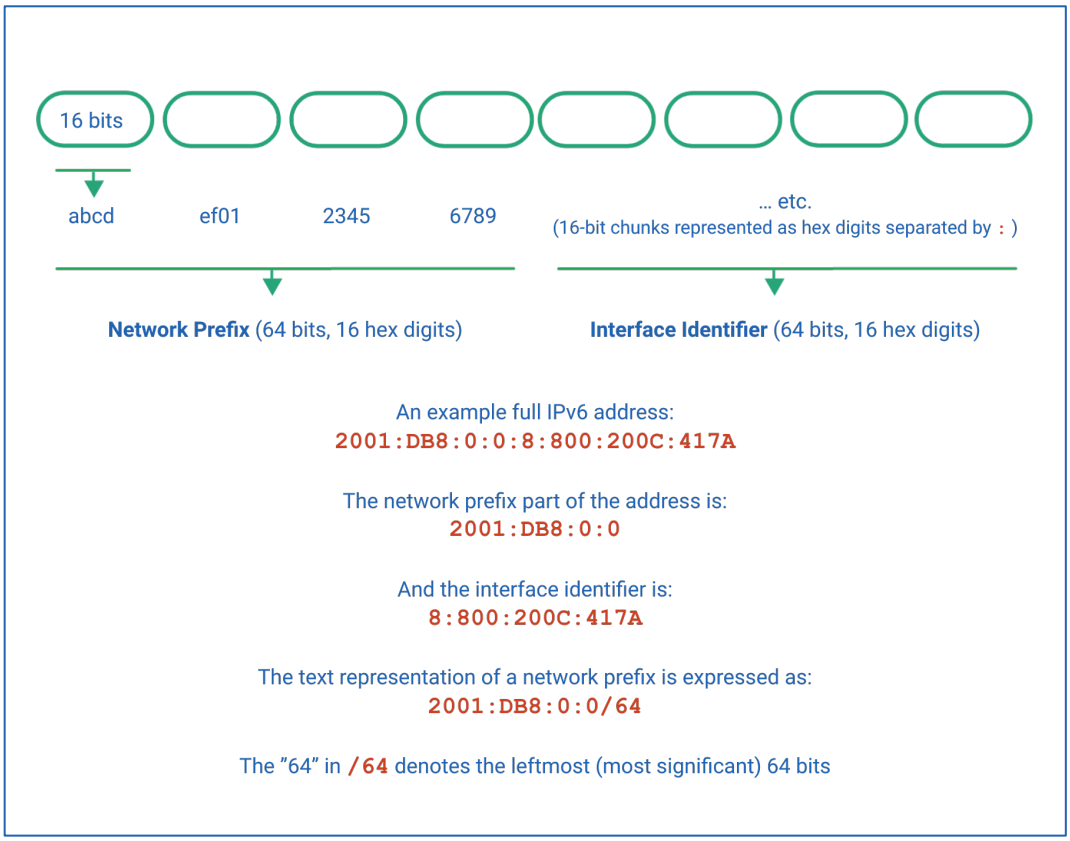

Diagram 1: Anatomy of an IPv6 Address

[5] https://www.ietf.org/

[6] https://datatracker.ietf.org/doc/html/rfc4291

Diagram 1 illustrates the **layout** of IPv6 addresses. In each address there is a **network prefix** part of variable length noted by **/x**, where **x** is the number of bits in the network prefix, counting from the left. For IPv4, it is common to have a 24-bit network prefix, noted as a **/24**, denoted by, for example: **10.7.13/24**. In IPv6 an address might have a 60-bit prefix noted as a **/60**, for example: **2001:0DB8:0:CD30:/60**.

Also note that in IPv6 the rightmost 64 bits are the **interface identifier**.

In IPv6, there are several methods used to allocate IPv6 addresses to an interface on a device. **RFC 7721**[7] discusses considerations in generating interface identifiers both for privacy and security. Following the considerations in the RFC will make devices that do not need a permanent address nearly impossible to find through scanning. Sadly, we have observed that sometimes these considerations are not carefully followed, and implementations remain vulnerable to scans.

In residential networks, it is common for the network to delegate a prefix to a residential customer premises equipment (CPE) device and for that device, in turn, to create multiple subnets within the residential network through further prefix delegations or to assign addresses to interfaces of individual devices.

Devices may also generate their own IP addresses on individual interfaces. We use the term "**address assignment**" throughout this paper to refer to the assignment of an IPv4 address or the delegation of an IPv6 prefix.

[7] https://datatracker.ietf.org/doc/html/rfc7721

Current Reported Practices in IPv6 Deployment That May Lead to Unnecessary Vulnerability

A number of researchers around the globe have been tackling the issue of finding IPv6 devices attached to the Internet for purposes of measuring the Internet[8].

Unlike the easily scanned IPv4 Internet, IPv6 is too large for exhaustive search techniques— it has been estimated that, even on the best of processors with the best of connectivity, a complete scan of the IPv6 address space could take on the order of 2x10^25 years.

The identification and publication of techniques for finding IPv6 hosts in the vast address space, therefore, has been a traditional source of research. In fact, some of the earliest research in IPv6 network reconnaissance was published in 2016 as an informational RFC in the IETF (**RFC 7707**)[9]. Much of the subsequent research takes the recommendations of that RFC as a starting point.

A number of observations have been made and techniques proposed with some promise of success in finding IPv6 devices, though never arriving at any kind of exhaustive enumeration.

The same techniques, of course, can be used by malicious actors.

Researchers recognize this, and, whenever their work finds examples of techniques and practices that could make end devices more conspicuous than would be desirable, they typically point them out. In some cases, they even report them to network operators in the form of potential vulnerabilities.

[8] See **References** below, in particular **Gasser et al.**, **Gao et al.**, and **De Coster et al.**

[9] https://datatracker.ietf.org/doc/html/rfc7707

Researchers have also found cases of IPv6 devices that are not secured with the same kind of software protections that IPv4 devices have. For example, whereas networks commonly set devices not to respond to certain kinds of queries on the IPv4 interface by policy (perhaps by not allowing SSH login attempts to those devices from external address spaces), they typically fail to set similar policies in IPv6, possibly because they are not aware that those devices could be reachable over IPv6 and/or that security practices must be duplicated on both IPv4 and IPv6 interfaces.

That adds to some other observed practices in the research about scanning and finding IPv6[10] hosts in the larger address space:

- **It is often too easy to leverage publicly available address information.** An example of this is the address of a host that is placed in the DNS. Sometimes, someone who is operating the network will give other hosts in the network nearby addresses.
- **Hosts may be numbered carelessly.** Although IPv6 recommends randomized interface identifiers (the 64-bit part of the address that identifies particular hosts[11]), implementers will simply number the rightmost bits in order (1, 2, 3) or a similar succession in a different part of the interface identifier that is easily found. We observe this in simple analysis of practice in networks where we have IPv6-enabled devices. Network operators and researchers also report observing this practice in customer networks.
- **Operators may forget to enable protections on the IPv6 connections that are effectively enabled on IPv4.** Not specifically having to do with addresses and scannability, this practice results in situations where, for example, an operator may have a policy not to accept Telnet connections and set a filtering rule to prevent that, but only in IPv4.

[10] Enabling IPv6 by default is an excellent practice to increase the uptake of IPv6. However not all operators of networks are aware that this has been done in an implementation and, thus, miss critical vulnerabilities.

[11] See **Terminology and Definitions** above.

Survey of IPv6 Address Assignment Practices

Many Internet service providers (ISPs) now enable IPv6 to home subscribers, often by default, so that a home user may unknowingly be using IPv6. IPv6 addresses are provided to any device that is IPv6 capable.

Our research showed that in one home network, for example, an IPv6 address is assigned to the printer, and the printer prefers to use IPv6. Ironically, the only reason one can use that printer is because it is IPv6-capable, and that is because, for some reason, the router connected to it no longer forwards IPv4 traffic! The printer uses a cryptographically generated 64-bit interface identifier, but it does not vary this interface address over time and, so, if this address were exposed to the public Internet, it could be reliably found later[12].

In the United States, GCA has some measured results from AT&T home Internet and T-Mobile 5G home Internet, and some limited measurements from Comcast residential Internet service. Analyzing those measurements is out of the scope of this piece of research but, still, they point to the fact that specific configurations may result in scanning vulnerabilities in IPv6-attached devices.

In our limited experience of operating IPv6-enabled devices in these domestic networks, we see a range of implementation and operation strategies for public IPv6 addresses. Thus, we typically see two or three IPv6 addresses configured on each interface, whether the device is a Linux-based device or a Mac[13]. On these types of devices, one of the addresses is always a cryptographically generated address. This address remains constant over long periods of time,

[12] A better practice would be to change the interface identifier part of the address periodically. The local protocols that locate the printer do not rely on prior knowledge of an IP address but discover the presence of the printer dynamically, so long-term persistence is unnecessary.

[13] We have very little practical experience with devices running Windows or any of the mobile-specific operating systems, so our comments will be based on observations of operational behavior of Linux and Mac devices over several months.

perhaps indefinitely. Interestingly, we do not see this address used in any outbound traffic, and it never appears to be reachable from devices in other networks.

A second address is generated with a random 64-bit interface identifier. This address is labeled as the "dynamic address." It is used for most, if not all, external communication over IPv6. This address will change at least every 24 hours, and sometimes more frequently (on Mac machines, but not Linux machines, it will change every time the system sleeps). This address, when known, is always reachable over IPv6 from an external address.

On some networks, both Linux and Mac machines generate a second dynamic IP address that has the same network prefix, but in which the interface identifier is all zeros except for the last 16 bits (and in some cases, the last 8 bits). This IP address, although labeled "dynamic," changes only very infrequently if at all, and apparently only with a system restart.

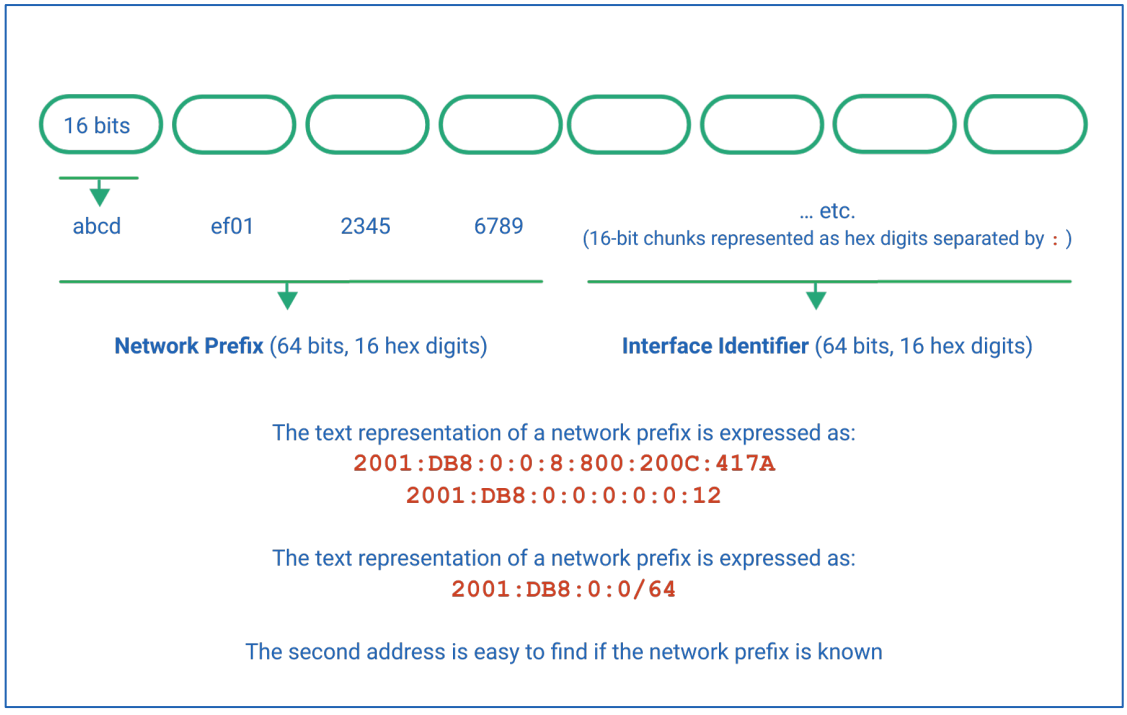

Diagram 2: Two Observed IPv6 Addresses

It is not clear what this second dynamic IP address is used for, and it does not respond to requests from outside the network or from requests within the same home network. Such an address would be somewhat easy to find at the time of this writing if the prefix assigned to the home network is known, regardless of whether that assignment is of a /56, a /60, or a /64, because the scan would be at most 24 bits (8 bits to determine the actual prefix if searching a /56 to find a /64, and 16 bits if the last 16 bits of the address are the only address fields that are populated— this is less than the number of bits to scan the entire IPv4 address space).

A landmark study[14] that analyzes data from both the RIPE Atlas network and observed data from a large Content Delivery Network (CDN) operator sheds some additional light on various practices of IPv6 address assignment.

The RIPE Atlas[15] network operates over 3000 dual-stacked probes in networks around the globe. These probes regularly report their IPv4 and IPv6 addresses allowing researchers to make observations about address assignment practices in dual-stack networks over long periods of time.

The CDN data provides a unique perspective based on content caches spread across the globe responding to data requests. Operator observations from this network are used to corroborate observations from the RIPE Atlas.

These observations and the comparisons between them allow a unique insight into evolving operational practices in a wide-ranging set of networks around the globe.

The study showed that, indeed, IPv6 prefix assignments to home networks are much longer lived in IPv6 than in IPv4. Such longer-lived prefixes may make it easier to target devices within

[14] See **Padmanahban, et. al.**, under **References**.

[15] https://atlas.ripe.net

a particular home or make the scanning easier if combined with sloppy interface identifier address assignments.

Exhaustion of the IPv4 address space leads operators to conserve IPv4 addresses, reclaiming addresses that are no longer used, and leading to frequent changes of IPv4 addresses. Many network operators have been deploying address sharing technologies such as carrier-grade-NAT that also lead to more frequent changes in addresses for individual subscribers. It is common for IPv4 addresses to be assigned to individual subscribers for as little as 24 hours, and frequently for intervals less than two weeks in many networks.

But in IPv6, the lack of incentive to conserve assigned prefixes leads to practices where a prefix is assigned to individual subscribers for months, as observed in the two data sets.

This IPv6 practice is interesting from a privacy point of view in that, regardless of how interface identifier assignment is done on devices, if a particular home network is assigned a network prefix and that prefix does not change, it becomes relatively easy to track users, at least within a specific home network, based on a long-lived prefix.

The study showed that some network operators renumber, that is, assign new network prefixes, quite frequently. German network operators, in particular, reassign prefixes sometimes on a daily basis. For most network operators assigning IPv6 prefixes, however, they assign prefixes to individual subscribers and do not change them for quite extensive periods of time.

Some Observations of IPv6 Scanning Activity

At the time of this writing, the amount of IPv6 scanning activity is minute compared to IPv4 scanning. This is, of course, mainly due to the comparative difficulty of scanning the IPv6 space in reasonable time. But also, research is limited, and the ability to discover vulnerable devices in IPv4 is much easier. One might expect this to increase when there is an increase in devices that are only reachable over IPv6.

One easy-to-check example is the public reports from Dataplane[16], in particular, of SSH scanning. Dataplane has a network of 300 sensors deployed in metropolitan areas on six different continents. These sensors are passive software entities that report activity directed toward them. Their sensors are not honeypots, so, for example, if they detect SSH login attempts, they do not allow the logins to succeed and then record the subsequent activity in a way that an SSH honeypot would. Nonetheless, their sensors do detect SSH scanning activity.

The organization collects the activity on their sensors and produces reports of various signals[17] of network activity that are free for non-commercial use. One of those signals is SSH login attempts[18] on their global networks of sensors.

That signal records attempts to make SSH logins at their sensors from IPv6 space every day. However, the amount of login attempts in IPv6 is tiny compared to IPv4— for instance, on one

[16] https://dataplane.org/

[17] Dataplane collects various kinds of activity directed toward the DNS, using IPv6, SIP activity, SMTP activity, SSH and Telnet login attempts, and requests to initiate VNC RFB sessions.

[18] The SSH login attempt signal is published at https://dataplane.org/signals/sshclient.txt. Their description of what is reported: "Entries below are records of source IP addresses that have been identified as performing SSH client protocol negotiations. Each entry is sorted according to a route originating ASN. An entry for the IP address is listed only once even if there are multiple origin AS (MOAS) announcements for the covering prefix. We use the PyASN IP address."

recent day, the Dataplane signal reported probes from approximately 36,000 distinct IP addresses, of which only 317 were IPv6[19].

On the other hand, Shadowserver[20] helpfully shed light on its practice of scanning for vulnerabilities in IPv6 networks in a report to the Forum of Incident Response and Security Teams (FIRST) conference in 2022[21].

For some years, Shadowserver has scanned the IPv4 repeatedly every day, producing reports[22] for national and regional government entities and for network operators. More recently, they began to offer vulnerability assessments in the IPv6 space as well, to quickly discover that techniques other than mass scanning were needed to identify vulnerable IPv6 devices.

At the FIRST conference in June 2022, they reported scanning known IPv6 addresses using curated lists from DNS AAAA data (passive DNS data), from the IPv6 hitlist repository described below[23], from certificate transparency streams, from sinkholes, and from partners. They scan

[19] Approximately half of the address blocks were in Hurricane Electric (2001:470:1), and half from Constantine Cybersecurity (2a06:4800:1000), with just a handful from Palo Alto Networks (2604:a940). Some abuse forums report that the addresses originated in Hurricane Electric are actual Shadowserver scanners. Shadowserver does not appear to have any of its own IPv6 address space, so this may be the case. Constantine Cybersecurity, for its part, is a business name for Internet-Measurement.com, which is operated by Driftnet (https://driftnet.io). Driftnet performs a number of different types of scans and reports the result of these scans as a service.

[20] The Shadowserver Foundation (https://www.shadowserver.org) is a not-for-profit company that provides a number of security related services. They operate a honeypot, a DDoS blackhole, and perform regular scans to generate reports of vulnerabilities in networks around the globe. They are also partners of GCA and members of the Nonprofit Cyber alliance (https://nonprofitcyber.org).

[21] See **De Coster, et. al.**, under **References**.

[22] Shadowserver currently documents 124 reports they provide to network operators: https://www.shadowserver.org/what-we-do/network-reporting/. Many of these reports are from various scans of the entire IPv4 address space conducted multiple times per day including reports from honeypots collecting attempted compromises over SSH and Telnet, reports of various open services (such as HTTP proxies), reports on whether various SQL services are open, on whether a host will accept QUIC connections, and many others.

[23] https://ipv6hitlist.github.io

these known addresses looking for vulnerabilities in SSL, SMTP, Telnet, SSH, HTTP, MySQL, FTP, and RDP. **Diagram 3** shows the nine services Shadowserver reported scanning using IPv6 (and how many responses they received on each service) when they announced their IPv6 scans[24].

Scan	Port	Responses
SSL	443/tcp	8 375 757
SSL	8443/tcp	230 687
SMTP	25/tcp	408 367
Telnet	23/tcp	24 377
SSH	22/tcp	849 193
HTTP	80/tcp	118 704 816
HTTP	8080/tcp	415 197
MySQL	3306/tcp	1 697 870
FTP	21/tcp	1 405 128
PostgreSQL	5432/tcp	10 712
RDP	3389/tcp	21 887

Diagram 3: Shadowserver IPv6 scans with reported responses in November 2022

Based on their scanning activity, they are able to identify particular IoT devices around the globe and to make observations about vulnerabilities. Some of the interesting findings from this work include the discovery that IPv6-capable devices use more up-to-date protocols, but also that IPv6 interfaces are much less likely to use security rules blocking scanning activity from outside.

Because of the difficulties of discovering devices through simple IPv6 scanning, a number of additional techniques have been developed.

[24] https://www.shadowserver.org/news/hello-ipv6-scanning-world/

Targeted scanning finds devices with known IPv6 addresses through public repositories such as the DNS and then makes limited scans in the address space near the public IP address, since, too often, administrative policies and simplicity of management involve the configuration of groups of devices with successive IP addresses.

Lists of IP addresses around which to scan are often developed by researchers and called "IPv6 Hitlists."[25] The largest of these lists are typically developed using private means, by analyzing server logs or developing ads that generate hits on sensors that track IPv6 sources. The next largest source of addresses to use seems to be public DNS records. There are also smaller lists of public resources, such as bitcoin IPv6 address lists, and even sources such as the public resources developed from the RIPE Atlas, as described above.

None of these are particularly expansive compared to simple scans of the IPv4 address space but they are the resources that third parties may use to find IPv6-enabled devices.

One less well-documented technique for finding IPv6 addresses for devices is through using specially constructed IPv6 packets (ICMP packets) in networks with known network prefixes. This technique was presented at Black Hat 2021[26]. However, it does not seem to be a generally usable technique, but one that relies on specific device vulnerabilities which seem likely to be fixed when understood.

[25] See **Murdock et al.**, under **References**.

[26] See **Gao et al.**, under **References**.

Existing IPv6 Address Assignment Policy Recommendations for Networks

The IETF and Réseaux IP Européens (RIPE)[27] have both published policy recommendations for network operators. The IETF publications reflect the recommendations of technologists from network equipment vendors, researchers, and network operators. The RIPE recommendations reflect the opinions of network operators in the RIPE region (Europe), but also includes perspectives from operators around the world who participate in the RIPE discussions.

RFC 7934 —also a best current practice (BCP), under **BCP 204**[28]— includes a thorough list of recommendations for address assignment for IPv6-enabled devices, including assigning multiple IPv6 addresses to devices using whatever address assignment technique is used operationally. It also recommends not restricting devices from configuring their own addresses on network interfaces and using those too. These are recommendations for network operators and equipment vendors around the number of addresses to be made available to a device.

Based on those recommendations, it might be possible to see a large number of individual IPv6 addresses on a single link between an IPv6-enabled device and a network.

[27] https://www.ripe.net/

[28] https://datatracker.ietf.org/doc/html/rfc7934

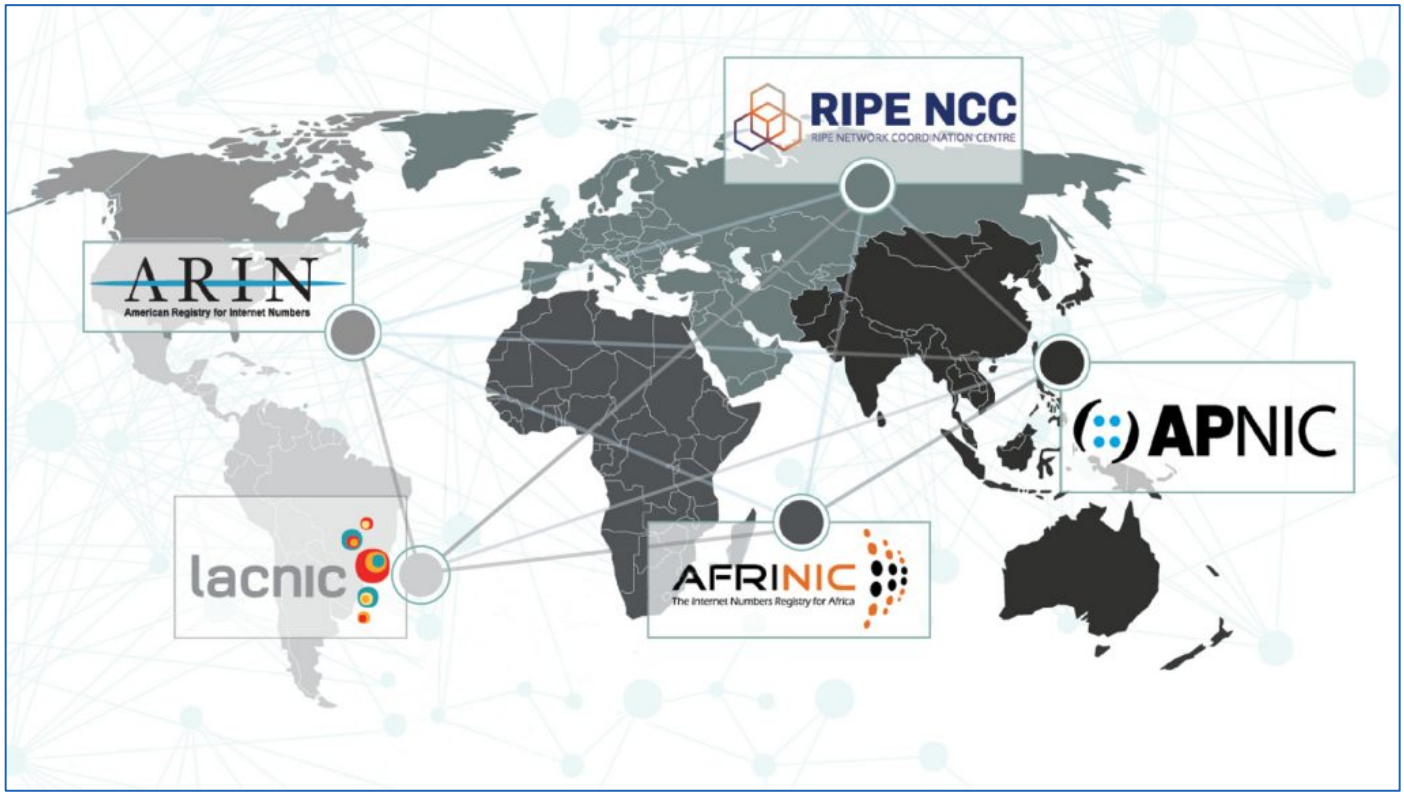

Diagram 4: The five Regional Internet Registries (RIRs) and their respective economies

There are five Regional Internet Registries (RIRs) that handle number resources in various regions across the globe. **Diagram 4** above shows the regional coverage of each. The registries often collect input from their communities of network operators and make policy recommendations about things such as IP address assignments and best operational practices.

The RIPE region has offered policy recommendations on IPv6 address assignment practices for network operators (**RIPE-690**)[29] that could be summarized as follows:

- The favored allocation of addresses from ISPs to all networks is to use a /48 prefix in such assignments

[29] https://www.ripe.net/publications/docs/ripe-690/

- A concession is made that some operators will choose to allocate a /48 for business customers and a /56 for residential customers
- There is a strong discouragement to assign anything longer than a /56 while pointing out that a /64 or longer is actually outside of specifications of the IETF (RFC 7934)
- In cellular networks, a /64 will be assigned to a phone for each connection, but there is a recommendation to use the previous recommendations on a cellular link providing data service in a home or business (5G hotspots are an example of this kind of operation)

Suggested IPv6 Best Practices at the Device Level

The study of address assignment practice at the network level sheds some interesting light on network-level address assignment policies. Based on those observations, existing recommendations for network operators from the IETF and the RIRs would improve the behavior of networks and devices.

This section, however, is focused on best practices at the device level.

It may be possible to provide relative security for individual devices regardless of the network level policy simply because of the large amount of address space available within the interface identifier part of the IPv6 address field.

Best practices about address assignment include the following:

- **Use random interface identifiers.** Some people configure interface identifiers from a small set, and often these are close to something that is stored in the DNS (clustered around a public address), whereas some others configure them from a really small and predictable set. For a device that does not need to be known in the DNS and does not need to use a long-lived address, picking an address for external communications where the interface identifier is generated randomly across all 64 bits and changed frequently will make the device nearly impossible to find.
- **Make frequent address changes** (every 24 hours is common), where the interface identifier does not need to be long-lived.
- **Make frequent prefix changes** (many German network operators apparently[30] do this every 24 hours).
- **Discard addresses that are no longer used.** In our research, for example, Linux laptops retain the public address it has used to reach a website even though the laptop has

[30] See **Padmanabhan et al.**, under **References**.

generated a new address now used for outbound traffic. That address remains reachable from external devices.

General security best practices:

- **Be aware when your network is forwarding IPv6 traffic**. This seems basic, but, apparently, it is often not obvious; this is an issue that points to the need for further training, as reported regularly in public operator group meetings.
- **Be aware when individual devices are IPv6 capable and are forwarding IPv6 traffic** ("On" by default).
- **Make sure that security policies are specified and implemented in both IPv4 and IPv6 networks.** There may be legitimate reasons to implement different policies in IPv4 and IPv6, but, in general, having identical policies is a more likely case. Note it may not be obvious that the policies need to be configured for both IPv4 and IPv6 at either the device level or the network or router level.

Areas of Further Exploration

Having observed existing practice in our own research of networks and in the published literature, we have identified a number of areas ripe for further research.

In particular, these areas of exploration may further facilitate the deployment of honeypots and honeyfarms that may attract IPv6 attacks and provide additional data for research on distribution of malicious software:

- **Find out who assigns these easy-to-find addresses.** Is it particular implementations of Linux[31]? Is it an operator policy? Is it router-based? If it is consistently one of these, there may be a direct action to update code or operational practice, and there might also be a policy recommendation —perhaps through the RIRs— to improve operational practices.
- **Rely on the route collectors used by many network operators** (and often located at IXPs). These systems collect routes with /56 and /64, and sometimes even longer prefixes. The longer a prefix is visible, the easier it is to find hosts within it. It is trivial to scan the IPv4 Internet entirely, so it is also trivial to scan a /96 in IPv6 space based on search size. Therefore, when there is a known prefix approaching this length, it becomes more easily scannable[32]. So the question becomes— how easy are these longer prefixes to obtain? How easy is it to find hosts on them?
- **Use GCA's developed ProxyPot, a proprietary traffic-capturing sensor.** This sensor allows for the creation of honeyfarms with similar characteristics to the ones available in the market, but also with the advantage of adjusting them to specific needs, such as adding detecting attacks against IPv6 interfaces.

[31] Linux and Linux-based operating systems are the typical OS for use in IoT devices. And as we observed earlier, many of the malware distribution networks now look not just for IoT devices, but for any system running Linux, so potentially large and powerful server implementations.

[32] See **Sediqi et al.**, under **References**.

Requirements to Extend ProxyPot to Capture IPv6 Attacks

The following section covers the basic requirements for extending development of GCA's ProxyPot technology to capture IPv6 attacks:

- Enabling IPv6 interfaces in ProxyPot to do a full analysis of attacks against honeypots with IPv6 enabled
- Enabling an IPv6 web proxy to attract traffic toward the existing honeypot before having a full IPv6 implementation in ProxyPot
- After enabling IPv6 in ProxyPot, making it attractive to attackers using some of these insights:
    - Putting an IPv6 address into the DNS, exposing the subnet that something is on (a simple web server perhaps)
    - Putting some emulated devices into the same subnet with addresses that are easily found (prefix ::0, for instance)
    - Using simple interface identifiers that never change
    - Being active over IPv6 to suspicious domains, for instance, by responding to phishing attempts where IPv6 is enabled (this may be too much of a stretch)
    - Performing comparative testing of honeypots that really are IPv6-only vs IPv4 and IPv6-enabled to observe differences in the attacks made on the IPv6 interface[33]

[33] This is particularly interesting because, in general, techniques for finding IPv4 and IPv6 devices may not find both interfaces on the same device, but techniques have been reported where people are able to coax address information out of a device on IPv6 interfaces once an IPv4 interface has been discovered.

References

The following research publications are some of the most relevant for this report:

- **Czyz, Jakub, and Matthew Luckie, Mark Allman, and Mark Bailey.** "Don't Forget to Lock the Back Door! A Characterization of IPv6 Network Security Policy." *Proceedings of the ISOC Network and Distributed System Security Symposium* (San Diego, United States, 2016). The Internet Society (ISOC). https://www.ndss-symposium.org/wp-content/uploads/2017/09/dont-forget-lock-back-door-characterization-ipv6-network-security-policy.pdf.
- **De Coster, Dave, and Piotr Kijewski.** "Internet Spelunking. IPv6 Scanning and Device Fingerprinting." FIRST Conference (Dublin, Ireland, 2022). Forum of Incident Response and Security Teams (FIRST). https://www.first.org/resources/papers/conf2022/08-InternetSpelunking-KijewskiandDeCoster.pdf.
- **Fukuda, Kensuke, and John Heidemann.** "Who Knocks at the IPv6 Door? Detecting IPv6 Scanning." *Proceedings of the 2018 Internet Measurement Conference, IMC '18* (Boston, United States, 2018). Association for Computing Machinery (ACM). https://dl.acm.org/doi/10.1145/3278532.3278553.
- **Gao, Shuping, Xingru Wu, and Jie Gao.** "New Ways of IPv6 Scanning." Black Hat Europe 2021 (London, United Kingdom, 2021). Black Hat. https://i.blackhat.com/EU-21/Wednesday/EU-21-Shupeng-New-Ways-of-IPV6-Scanning.pdf.
- **Gasser, Oliver, Quirin Scheitle, Pawel Foremski, Qasim Lone, Maciej Korczyński, Stephen D. Strowes, Luuk Hendriks, and Georg Carle.** "Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists." *Proceedings of the 2018 Internet Measurement Conference, IMC '18* (Boston, United States, 2018). Association for Computing Machinery (ACM). https://dl.acm.org/doi/10.1145/3278532.3278564.
- **Murdock, Austin, Frank Li, Paul Bramsen, Zakir Durumeric, and Vern Paxson.** "Target Generation for Internet-wide IPv6 Scanning." *Proceedings of the 2017 Internet Measurement Conference, IMC '17* (London, United Kingdom, 2017). Association for Computing Machinery (ACM). https://dl.acm.org/doi/10.1145/3131365.3131405.

- **Padmanabhan, Ramakrishna, John P. Rula, Philipp Richter, Stephen D. Strowes, and Alberto Dainotti.** "DynamIPs: Analyzing Address Assignment Practices in IPv4 and IPv6." 16th International Conference on Emerging Networking Experiments and Technologies, CoNEXT '20 (Barcelona, Spain, 2020). Association for Computing Machinery (ACM). https://dl.acm.org/doi/10.1145/3386367.3431314.
- **Richter, Philipp, Oliver Gasser, and Arthur Berger.** "Illuminating Large-Scale IPv6 Scanning in the Internet." *Proceedings of the 2022 Internet Measurement Conference, IMC '22* (Nice, France, 2022). Association for Computing Machinery (ACM). https://dl.acm.org/doi/abs/10.1145/3517745.3561452.
- **Sediqi, Khwaja Zubair, Lars Prehn, and Oliver Gasser.** "Hyper-specific Prefixes: Gotta Enjoy the Little Things in Interdomain Routing." RIPE 86 Conference (Rotterdam, Netherlands, 2023). Réseaux IP Européens (RIPE). https://www.ripe.net/media/documents/11-11-hsp_presentation_for_ripe_20230522.pdf.
- **Zhong, Jie, and Xiangning Chen.** "Research on DDoS Attacks in IPv6." *Proceedings of the 4th International Conference on Computer Science and Application Engineering* (2020). https://dl.acm.org/doi/10.1145/3424978.3425020.

The following best current practices (BCP) and requests for comments (RFC) have also been referenced across this report:

- **[RFC 4291]** Hinden, Robert and Steve Deering. "IP Version 6 Addressing Architecture." February 2006. https://datatracker.ietf.org/doc/html/rfc4291.
- **[RFC 7278]** Byrne, Cameron, Dan Drown, and Ales Vizdal. "Extending an IPv6 /64 Prefix from a Third Generation Partnership Project (3GPP) Mobile Interface to a LAN Link." June 2014. https://datatracker.ietf.org/doc/html/rfc7278.
- **[RFC 7707]** Gont, Fernando, and Tim Chown. "Network Reconnaissance in IPv6 Networks." March 2016. https://datatracker.ietf.org/doc/html/rfc7707.

- **[RFC 7721]** Cooper, Alissa, Fernando Gont, and Dave Thaler. "Security and Privacy Considerations for IPv6 Address Generation Mechanisms." March 2016. https://datatracker.ietf.org/doc/html/rfc7721.
- **[RFC 7934] [BCOP 204]** Colitti, Lorenzo, Dr. Vinton G. Cerf, Stuart Cheshire, and David Schinazi. "Host Address Availability Recommendations." July 2016. https://datatracker.ietf.org/doc/html/rfc7934.
- **[RIPE-690]** Žorž, Jan, Sander Steffann, Primož Dražumerič, Mark Townsley, Andrew Alston, Gert Doering, Jordi Palet, Jen Linkova, Luis Balbinot, Kevin Meynell, and Lee Howard. "Best Current Operational Practice for Operators: IPv6 Prefix Assignment for End-Users: Persistent vs. Non-Persistent, and What Size to Choose." October 2017. Réseaux IP Européens (RIPE). https://www.ripe.net/publications/docs/ripe-690/.

Appendix: A Major Operator's Perspective on Actual Practices of IPv6 Prefix Delegation

In this appendix, we discuss IPv6 address allocation practices from the perspective of a major operator who operates many different kinds of services, including mobile and fixed broadband, corporate network customers, and hosting of various kinds of services.

In a network with such a wide diversity of operations, IPv4 address allocation practices are complex. The complexity of address allocation and assignment is due to the operational realities of diverse service offerings, the history of network amalgamation, the reality of a limited supply of IPv4 addresses, the opportunity to consolidate IPv4 address usage because the IPv4 addresses in use are a valuable and saleable asset, and needs to adapt practices in order to reduce operational complexity.

The operational reality of diverse service offerings includes offerings to both large and small businesses and the kind of networks they operate, and the diversity of equipment they use.

Large networks are often operationally an amalgamation of networks designed by different network architects who have diverse approaches to manage devices and addresses within a network. Sometimes these networks are owned by the same company initially but often networks are brought together having originated in different companies with quite different business and operational models.

Although there is a desire to consolidate these operationally as well as organizationally, many times it is easier to continue to operate in their original configurations as resources to consolidate and normalize are often at a premium.

Of course, different kinds of network technologies often come with unique practices, not only in address allocation and management but also in all sorts of operational realities (e.g. security, routing).

Finally, even within the same network deployment of different kinds of equipment with different kinds of capabilities, equipment failure and replacement, and asynchronous upgrade cycles can lead to diverse operational practices. Also, both in business and residential settings, customers often bring their own devices with diverse capabilities to the network.

While one might hope to have something of a clean start with the deployment of a new technology such as IPv6 —or at least an opportunity to consolidate around a smaller set of operational practices—, this is often more difficult than one might hope.

One of the main complicating factors in some settings is when the ISP provides Internet services, but the customer of the ISP supplies the home router, which is a crucial part of address assignment within the home.

The variety of capabilities of such equipment makes it difficult to determine a policy that is both generous and consistent, so, in practice, the ISPs provide a more limited IPv6 allocation in such situations, often just a /64 to the premises, whereas in the case of an ISP-supplied CPE device, a /60 or a /56 is supplied.

For small businesses in a similar network configuration, it is often possible to assign a /56 to such a network because the customer premises device is designed for a small business network and has more sophisticated address management capabilities.

In all these settings, the assignment of addresses beyond the first hop customer premises equipment is up to that equipment or to devices behind it. When the customer premises equipment is supplied by the ISP, the ISP is often able to specify, test, and guarantee behavior that is more sophisticated than what otherwise might be available.

For simpler service offerings, such as a dedicated home video device, a /64 is often deemed sufficient.

In the mobile data side of network operations things are, in fact, often quite clean. Several operators report complying with the recommendations in RFC 7278 in deployments of wireless home routers on a mobile network interface. Regardless of recommendations to assign a /48 or /56 for this kind of service, assigning a /64 is common in these types of networks.

New York
31 Tech Valley Drive
East Greenbush
NY 12061
UNITED STATES

London
City of London Police
3rd Floor
Guildhall Yard East
London
EC2V 5AE
UNITED KINGDOM

Brussels
Scotland House
City Office in Brussels
(c/o Global Cyber Alliance)
Rond Point Schuman 6
1040 Brussels
BELGIUM

www.globalcyberalliance.org

GLOBAL CYBER ALLIANCE