Information Visualization

Accommodating IPv6 Addresses in Security Visualization Tools[*]

David Barrera and Paul C van Oorschot

The online version of this article can be found at:

Published by:

$SAGE

Additional services and information for *Information Visualization* can be found at:

Email Alerts: http://ivi.sagepub.com/cgi/alerts

Subscriptions: http://ivi.sagepub.com/subscriptions

Reprints: http://www.sagepub.com/journalsReprints.nav

Permissions: http://www.sagepub.com/journalsPermissions.nav

Citations: http://ivi.sagepub.com/content/10/2/107.refs.html

>> Version of Record - Apr 1, 2011

What is This?

Downloaded from ivi.sagepub.com at CARLETON UNIV on June 4, 2012

Original Article

Accommodating IPv6 addresses in security visualization tools☆

David Barrera* and
Paul C. van Oorschot

School of Computer Science, Carleton University, 1125 Colonel By Drive, 5302 Herzberg Building, Ottawa, Ontario, Canada K1S 5B6.
E-mail: dbarrera@ccsl.carleton.ca

*Corresponding author.

Abstract Visualization is used by security analysts to help detect patterns and trends in large volumes of network traffic data. With IPv6 slowly being deployed around the world, network intruders are beginning to adapt their tools and techniques to work over IPv6 (versus IPv4). Many tools for visualizing network activity, while useful for detecting large-scale attacks and network behavior anomalies, still only support IPv4. In this article, we explore the current state of IPv6 support in some popular security visualization tools and identify the roadblocks preventing those tools from supporting the new protocol. We propose a filtering technique that helps reduce the occlusion of IPv6 sources on graphs and enables IPv4 visualization tools to display both IPv4 and IPv6 sources on a single graph. We also suggest using treemaps for visually representing the vast space of remote addresses in IPv6.
Information Visualization (2011) **10**, 107–116. doi:10.1057/ivs.2010.9; published online 11 November 2010

Introduction

IPv4, the well-known protocol employed for practically every type of network communication, has been used for over 25 years. Unfortunately, its limited address space is projected to be exhausted by August 2012.[2] With this date fast approaching, system administrators, vendors, ISPs and end users around the world are slowly adopting IPv6. As IPv6 becomes more widely deployed, adversaries are updating their tools to use it and are benefiting from the fact that IPv6 is often overlooked, disabled or otherwise unsupported in many network analysis tools. Although Internet-scale enumeration of IPv6 hosts is unfeasible, attackers are using advanced features in the new protocol for local network enumeration and reconnaissance.

In security visualization, remote (source) IP addresses are frequently used as one of the displayed fields to discover trends or patterns of activity in groups of hosts. As IP addresses uniquely identify devices on a network, visualizing their behavior provides better insight into the structure and location of networks. Many security visualization tools[3–5] have been designed to support IPv4 exclusively, for example, graphs that enumerate the 32-bit address space (or segments thereof) on axes. Working with the 128-bit addresses of IPv6 is not as simple.

In this article, we review why visualization of remote IP addresses is important, and why IPv6 remote addresses specifically are difficult to graph. We analyze the structure of IPv6 addresses as well as the current IPv6 address allocation trends and use our observations to develop three new techniques for visualizing the IPv6 address space. One technique filters out unused IPv6 address ranges, allowing analysts to visualize the effective IPv6 address space seen generating traffic during a network capture.

The second technique allows IPv4-only tools to display dual stack network captures. In the third technique, we propose the use of an existing visualization tool (treemaps) to visualize IPv6 data sets.

The rest of this article is structured as follows. Section 'Background and related work' reviews technical details of IPv6 and how it differs from IPv4, and survey-related work. Section 'Considerations for visualizing IPv6 data sets' considers the general architecture of visualization systems and issues that must be addressed to support IPv6. Section 'Data sets' describes the data sets used to demonstrate visualizations in this note. Section 'Visualization of remote IP addresses' proposes three new techniques for visualizing network traffic containing IPv6 addresses. Section 'Conclusion and future work' presents concluding remarks.

Background and Related Work

To better understand the security visualization challenges faced when trying to visualize IPv6, we first review how it differs from IPv4.

IPv6 review

IPv6 extends IPv4 from 32- to 128-bit addresses[6] and includes simplified headers that help improve packet processing performance. Although the IPv6 specification has existed for over 10 years, deployment has been relatively slow. An October 2008 study by Google[7] found that IPv6 penetration worldwide is still less than 1 per cent in any country (measured by the number of hosts requesting AAAA DNS records from Google domains). The study also found that 67.9 per cent of IPv6 users worldwide are using 6 to 4 (a transition mechanism that encapsulates IPv6 packets inside IPv4 packets to help connect 'islands' of IPv6 networks) as opposed to native, IPv6-only connectivity. For the United States and Canada, 95 per cent of IPv6 users are using 6 to 4, presumably because of the small number of ISPs offering native IPv6 in these countries. China and France, on the other hand, are over 70 per cent native.

It is important to note that low current use of IPv6 on the public Internet does not mean hosts using the protocol should be ignored. Despite the aforementioned statistics, IPv6 penetration is growing every day and is now enabled by default in many network applications. Hosts are capable of simultaneously using IPv6 for internal network communication, and IPv4 for Internet communication. Thus, even if a host does not have Internet-wide IPv6 connectivity, it might still be exploitable over IPv6 within its subnet.

Transition Mechanisms
To help speed up IPv6 deployment and help ease migration, several RFCs provide methods allowing IPv4 and IPv6 to inter-operate during an initial transitional phase.

These transition mechanisms bring new issues to security analysis. For instance, if a network application uses one of the transition mechanisms such as a tunnel (a point-to-point IPv4 link between an end-user and an ISP through which IPv6 packets are encapsulated) or 6 to 4, the source and destination IP addresses on those packets may not represent the actual communication endpoints.

Address Representation
IPv6 addresses are not written in the IPv4 dotted-decimal notation, but rather as a group of eight 16-bit *hextets* (as opposed to four 8-bit pieces in IPv4). Leading zeroes in each hextet can be omitted, and up to one group of two or more hextets consisting of only zeroes can be expressed as '::'. The following examples show different possible notations for an IPv6 address.

```
2001:0DB8:0000:0078:9ABC:0000:0000:0000
2001:0DB8:0:0078:9ABC:0:0:0
2001:DB8:0:78:9ABC::
```

The representation of prefixes in IPv6 is similar to the Classless Inter-Domain Routing notation in IPv4. The prefix length, which specifies how many of the leftmost contiguous bits compromise the prefix, is appended after the IP address. For example, a network block where 64 bits are used for assigning to hosts and 64 remaining constant would be expressed as 2001:DB8:0:78::/64.

Other differences from IPv4 and visualization implications

The updated IPv6 packet header omits rarely used fields and has been simplified to allow faster processing by network devices. As shown in Figure 1, headers now contain fewer fields, by eliminating the *header length, identification (Internet Protocol Identifier (IPID)), flags, fragment offset* and *header checksum*. The *version, traffic class, payload length, next header* and *hop limit* fields in IPv6 headers are either the same as in IPv4 or used for similar purposes, making them somewhat compatible in both analysis and visualization tools. The *flow label* field is unique to IPv6 and provides a method of tagging all packets belonging to the same end-to-end conversation.

Since IPv6 and IPv4 operate at the Internet layer of the IP stack, there are only differences between protocols at this specific layer. Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and any other protocols that run over IP are unchanged, and therefore visualization tools that graph only transport or application layer data generally do not need to be modified to work with IPv6. On the other hand, security visualization tools that display Internet layer data, such as IP addresses or other header fields, might require changes to work with IPv6.

The address space in IPv6 is far too large for brute-force enumeration to be feasible (attackers cannot sequentially probe all IP addresses in IPv6). This means that remote host discovery is much less likely (unless it targets specific, small subnets). However, there are some related

	0 - 3	4 - 7	8 - 15	16 - 18	19 - 31
0	Version	Header Length	Type of Service (ToS)	Total Length	
32	Identification (ID)			Flags	Fragment Offset
64	Time to Live TTL		Protocol	Header Checksum	
96	Source IP Address				
128	Destination IP Address				
160					
192					
224					
256					**IPv4**

	0 - 3	4 - 11	12 - 15	16 - 23	24 - 31
0	Version	Traffic Class	Flow Label		
32	Payload Length			Next Header	Hop Limit
64					
96	Source IP Address				
128					
160					
192					
224	Destination IP Address				
256					
288					**IPv6**

Figure 1: Packet Headers in IPv4 and IPv6.

IPv6 features, such as neighbor discovery and multicast, that should not be overlooked. Multicast addresses can allow a single device to discover all other IPv6 devices on a local network as described by Moore[8] using *The IPv6 Attack Toolkit*[9] (a collection of network tools for discovering and attacking IPv6-enabled devices).

Another interesting feature of IPv6 is address auto-configuration. Hosts can now assign themselves an IP address without the need for a Dynamic Host Configuration Protocol (DHCP) server. Using ICMPv6, a server broadcasts the IPv6 prefix which all hosts should prepend to their auto-configured address (this is designed to be collision free since the host's Media Access Control (MAC) address is used as part of the address). While address auto-configuration helps for network renumbering, it makes monitoring hosts on an IPv6 network a more complicated task.

Focus on remote IP addresses

IP addresses uniquely identify hosts on networks. The vast majority of network analysis tools use IP addresses either to group or classify hosts at the network level. Visualizing IP addresses is useful for detecting similar behavior among groups of hosts as well as mapping the origin of a connection to an approximate geographical location.

It is useful to distinguish between externally initiated connections and those that are internally initiated. In this article, we focus on the former; the remote hosts in question are *remote addresses*, while those assigned to devices located within a local network perimeter are *local addresses*. Remote and local addresses are also referred to as *source* and *destination* addresses, respectively.

A network administrator should always have local knowledge of the network being defended, that is, which segments of the address space are populated or active. For this reason, monitoring target or local IP addresses in IPv6 is somewhat easier than doing the same for remote addresses. For instance, even if an administrator has 2^{64} addresses available for assignment, in all likelihood, only a very small (known) subset of these will be actually used, and the corresponding devices can easily be monitored. Many security visualization tools[5,10] graph the host segment of the IP address for displaying internal addresses. With this approach, no information is lost from the IP address because the network segment of the address is constant.

Remote IP addresses, on the other hand, are more difficult to monitor, since it is impossible to know ahead of time which remote hosts will initiate connections, and visualizing a small portion of the address could hide information from the user. It is common to see tools visualizing the entire remote IP address or the network segment of the address, with the drawback of losing the host-level accuracy. In this article, we focus on visualizing remote IP addresses and attempt to limit the loss of accuracy.

Related work

We are not aware of any literature that specifically addresses visualization of IPv6 traffic. The existing

literature on security visualization also does not explicitly state that the focus is IPv4 traffic, but typically implicitly assumes that IPv4 is the underlying protocol of the analyzed network traffic. Several visualization tools[5,4,10] that provide good insight into remote IP behavior are limited to IPv4 addresses.

Nakamae et al[11] attempted to visualize the structure of an IPv6 network. This differs from our approach in several ways. They did not display IPv6 addresses in their visualizations or analyze IPv6 network traffic, but rather displayed the connections linking different ISPs (obtained from routing table entries) within a specific region. They do not aim to help visualize the behavior of groups of IPv6 hosts. They focus on representing the IPv6 address hierarchy (see Section 'Using IPv6 address hierarchy with treemaps') through an interactive 3D display. Their results produce attractive 3D visualizations rich in information (compared to 2D), with the accompanying trade-off of requiring the user to actively interact with the tool to see information that may be occluded by other parts of the graph. The Cooperative Association for Internet Data Analysis (CAIDA) has also produced visualizations displaying IPv6 autonomous system interconnections.[12]

Considerations for Visualizing IPv6 Data Sets

Most visualization tools follow these steps to visually represent network traffic captures:

1. *Capture raw data*: data off the wire is captured by a common network monitoring tool (Tcpdump, Wireshark, NetFlow and so on).
2. *Parse*: the raw dumps are then parsed to extract information from the data structures used by the monitoring tools. This extracted data is written to text files.
3. *Process and reformat*: text files are then processed for field selection and modified into the format required by the visualization tool.
4. *Display*: graphically represent the data.

Data capturing and parsing

Popular data *capturing* tools have included support for IPv6 for a few years. These tools are capable of detecting IPv6 in IPv4 encapsulation, detecting the different types of addresses and understanding the new security enhancements (for example, IPSec) that are built-in to IPv6. For example, the SiLK tools[13] allow the analyst to display IPv4 flows as IPv6 flows containing 6 to 4 addresses. SiLK is also being improved to support faster manipulation of IPv6 addresses using high-performance data structures.[14] As there is very good support for capturing IPv6, this part of the visualization process requires practically no changes. The main point is to make sure that network analysts enable the IPv6 capture feature as it might be disabled by default in some capture tools.

Once data have been saved to a capture file, parsers (generally written in scripting languages, for example, Perl or Python) are used to read relevant fields from the capture. Parsers work with plain text, so they should be able to extract IPv4 and IPv6 without much difficulty. When working with parsers, just as with capture tools, analysts need to make sure the parser is not discarding IPv6 records, since IPv6 addresses are written differently and may cause a poorly written IPv4-only parser to crash or ignore that part of the file.

Data processing and visualizing

After parsing the file, data for the visualization must be selected. Common network fields that are used in visualization are: source and destination address, the IPID and source and destination ports. Of these, only the source and destination ports remain unchanged in IPv6. Source and destination addresses have become 128 bits, and the IPID is removed from the header. Visualization tools must not try to extract an IPID (or other non-existent fields) from an IPv6 packet, nor attempt to display an IPv6 address when processing data records from an IPv4 space; existing tools that do so must be modified to prevent them from breaking when processing IPv6 traffic.

Current visualization tools generally do not support IPv6 in the sense that they assume the capture is IPv4 only, and therefore assume 32-bit addresses. For example, the Inetvis[3] cube is hard-coded for IPv4. If a dual stack capture (that is, containing both IPv4 and IPv6 data) is loaded, the tool will just ignore the IPv6 packets and not display them. NVisionIP[4] and Flovis[10] are other examples of tools where only 32-bit source addresses are currently supported. Koike et al[5] also describe a technique similar to NVisionIP in which source and destination addresses are visualized in a 2D matrix; only IPv4 addresses can be displayed.

Data Set

Our sample visualizations were generated from two distinct data sets, one containing only IPv6 traffic, and the other containing both IPv4 and IPv6 traffic (sometimes referred to as dual stack). The data sets were obtained from the MAWI Working Group Traffic Archive,[15] which is host to a large number of up-to-date traffic captures from several Internet backbone links.

IPv6-only data set

For the visualizations of Sections 'Whitespace filtering' and 'Using IPv6 address hierarchy with treemaps', we used a data set file from the MAWI Working Group's *Samplepoint-C* (one of their 16 sampling points), which contains 2 million IPv6-only packets, generated on 18 June 2008 between 9:00 and 13:36 and were anonymized

using a modified version of *tcpdpriv* to comply with WIDE standards. Payloads were removed, but protocols and ports were left unchanged by *tcpdpriv*. The IP addresses were scrambled, with sources sharing the same address prefix preserving this property (two IP addresses A and B on the same subnet X would both be placed in a new subnet Y) throughout the duration of the capture period. The data set contains 1924 (out of which 1466 generated only TCP and UDP traffic) unique source IP addresses generating traffic at an average 221 kbps. Traffic is distributed as TCP (50.3 per cent), UDP (29.6 per cent), and Internet Control Message Protocol (ICMP) (19.7 per cent) as well as a small number of other protocols.

Dual-stack data set

The dual stack data set used in Section 'Whitespace filtering' is a capture file from the MAWI Working Group's *Samplepoint-F*, which contains about 11.5 million packets captured over a 15-minute window on 3 January 2008. This data set in particular recorded a considerably higher rate of traffic (averaging 49.95 Mbps) than the IPv6-only data set over a relatively small period of time. Network traffic was also anonymized using the same process as the data set of Section 'IPv6 – only data set'. Most of the recorded packets are IPv4 (99.79 per cent), with 354 472 unique hosts, and the rest (0.21 per cent) are IPv6, with 3029 unique hosts. Traffic for this data set is distributed as TCP (63.59 per cent), UDP (28.66 per cent), ICMP (7.01 per cent) and a small number of other protocols.

Visualization of Remote IP Addresses

To represent IP addresses in IPv6, trying to grow the range of an axis from 2^{32} to 2^{128} does not work. One must

rethink how to display information to gain insight on how connections are distributed and where they come from. In addition, we seek to work with the existing visualization tools with as little disruption as possible, even though such tools might not be ready for IPv6. We provide three ideas in this direction.

Whitespace filtering

One of the problems for visualization in IPv6 is that the address space is very sparsely populated. Most of the left-most 4 bits are reserved by the Internet Engineering Task Force (IETF) for future use, which leaves only one prefix (2000::/3) that can be globally routed. To demonstrate how unevenly distributed the remaining addresses are, in Figure 2 (left), we plot the data set of Section 'IPv6 – only data set' using the full IPv6 address range on the *y* axis. Note that we can only distinguish six horizontal lines (corresponding to addresses beginning with 20, 24, 26, 2A, 3F and FE). Figure 2 (right) zooms into the bottom line of Figure 2 (left); what appeared to be a single address is actually a group of hundreds of addresses (displayed as horizontal lines). Even on a high-resolution display (for example, 1920 × 1200), we could potentially still have thousands of remote IP addresses overlapped onto a single row of pixels of the screen (not the case in this example, as the capture contained only 1924 source addresses).

In IPv4, for an analyst to detect patterns and trends using a visualization tool, it is useful to look at neigh-boring addresses (in the same subnet or class). This can help identify hosts that are collaborating or exhibiting similar behavior. Furthermore, being able to see the behavior of all subnets simultaneously proves to be useful in detecting Internet-wide attacks. IPv6 does not immediately add more hosts to visualizations, just a much wider
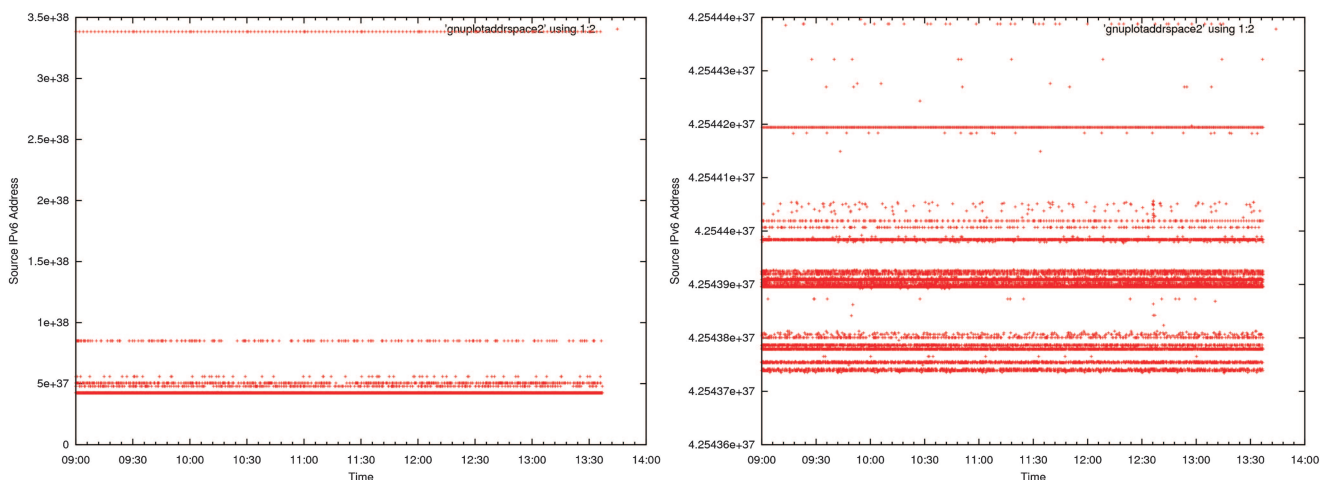
Figure 2: These graphs display the data set from Section 'IPv6 – only data set'. The *x*-axis displays a 4.5-hour window and the *y*-axis shows the source IPv6 address. The graph on the left shows hosts seen generating packets during the capture pereiod. The graph on the right is a 70 000× zoom on the *y*-axis of the left image focusing on the bottom solid line.
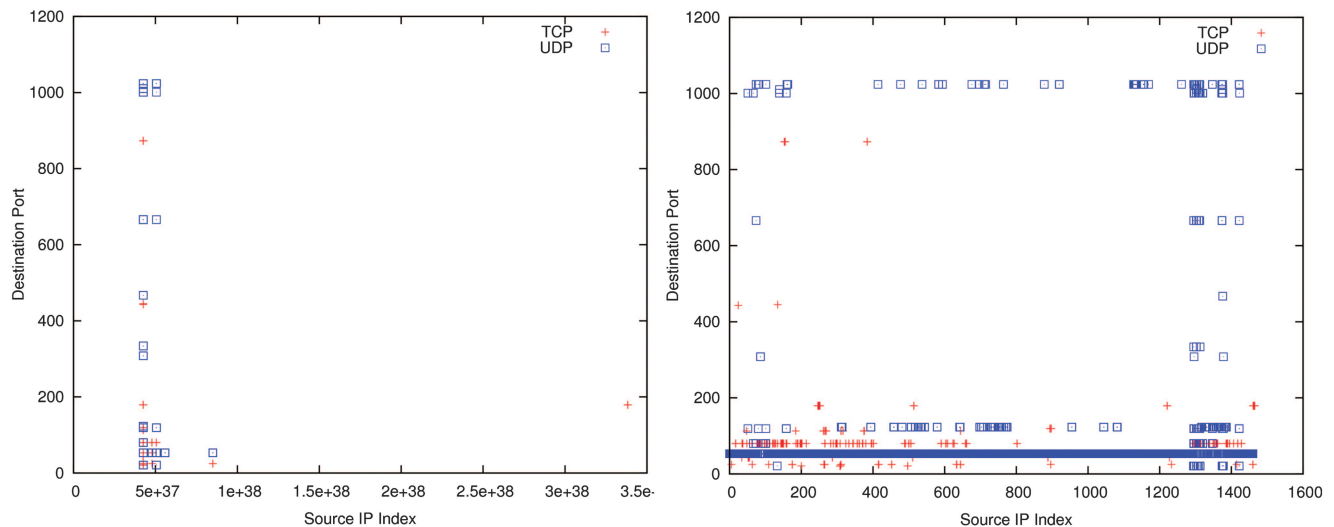
Figure 3: These graphs display the data set of Section 'IPv6 – only data set' (destination ports contacted by each source IPv6 host) before (left) and after (right) indexing the source IPv6 address and removing unused address space. The graph on the right uses an index value in place of the absolute IP address.

address range. Removing the unused or unpopulated addresses, that is, the whitespace (sometimes referred to as *darkspace* in scan detection literature), should restore the same, Internet-wide visualizations as with IPv4.

Eliminating the whitespace is easily done as follows. During the network capture, build a list of unique remote IPv6 addresses observed. To preserve relative placement (so that if address A numerically follows B, this relation is retained in the new list) of addresses, sort the resulting list by numerical value. Finally, when displaying data, display the index of the IP address in the list instead of the full address.

As an optional step, subnet information can be preserved in the graphs, by deliberately inserting gaps. We insert two null elements into the previously mentioned IPv6 address list between each populated 64-bit subnet. The end result is a list where all addresses in the same /64 subnet are contiguous, and all /64 subnets are separated by a small space. This optional step may help the user identify subnets visually rather than trying to otherwise validate findings in the original raw data set.

To save the analyst from changing his workflow by having to look-up the address that matches a given index, visualization tools could be modified to display the address when mousing over an index.

Figure 3 (left) presents our first data set as an example IPv6-only capture being visualized with no filtering at all. The *x* axis displays the 128-bit addresses as an integer value. Owing to the uneven distribution of addresses, as well as the huge IPv6 address space, the populated addresses all overlap on the left side of the scatterplot. In Figure 3 (right), unused address space has been eliminated, and each remote IPv6 address has been indexed. The resulting scatterplot is able to

display over 1400 remote addresses and what destination ports each one targets during the capture period. This filtered scatterplot reveals that traffic to UDP port 53 (Domain Name System (DNS)) was frequent throughout the entire capture period, and that no single host attempted connections to a large number of destination ports (noted by the absence of vertical lines in Figure 3 (right)). Note that the horizontal axis in Figure 3 (right) displays the *address index* rather than the full address.

Mapping IPv6 into unused IPv4 address space

The technique of Section 'Whitespace filtering' works exclusively with IPv6 data sets. As IPv6 and IPv4 can co-exist on a network, it is possible for a network capture to include both IPv6 and IPv4 packets. Furthermore, while migration from IPv4 to IPv6 takes place, it is likely that many networks will see both types of traffic. This section presents a technique that builds upon that of Section 'Whitespace filtering' to allow dual stack data sets to be displayed on IPv4-only visualizations (or limited to a tool assuming a 32-bit address space for display).

Owing to the currently low deployment of IPv6, for a given dual stack network, it is reasonable to assume that there will be more IPv4 traffic than IPv6 recorded at the network border. This may vary if traffic capturing is done internally (where most traffic could conceivably be IPv6), but that scenario is beyond the scope of this article. Mapping the currently small amount of IPv6 traffic into unused (that is, reserved for private or future use) portions of the IPv4 address space provides a visualization strategy that is capable of displaying both IPv6 and IPv4
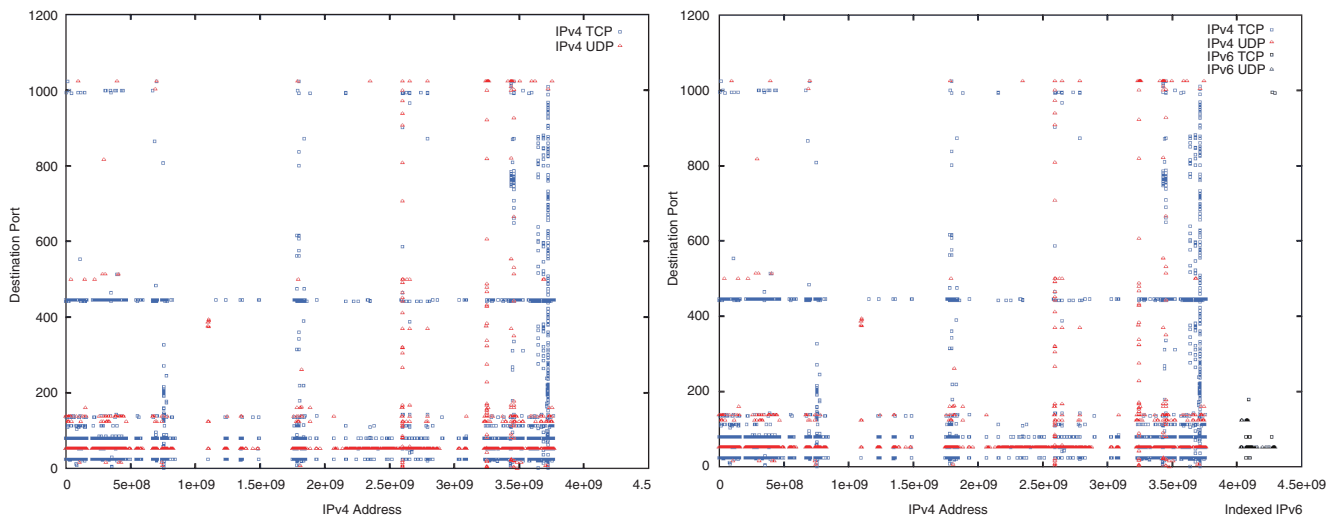
Figure 4: The graph on the left plots destination ports contacted only by IPv4 hosts in the dual-stack data set of Section 'Dual-stack data set'. The graph on the right maps IPv6 hosts into the unused portion of the graph, displaying the full data set.

traffic, while maintaining compatibility with current IPv4-only tools.

The mapping technique begins with indexing the data set, assigning a sequential, unique ID to each IPv6 address seen during the capture period (as described in Section 'Whitespace filtering'). IPv4 addresses are left unmodified during this step. Next, an unused IPv4 address range is selected. For this step, good candidate networks include those specified as private in RFC 1918,[16] or the entire Class E (range 240.0.0.0–255.255.255.255), which is currently reserved for future use.[17] Using, for example, RFC 1918's 10.∗.∗.∗ network will provide three full octets (2^{24} addresses) for mapping in IPv6 addresses, allowing up to 16 777 216 IPv6 addresses to fit in that range without overlap. Similarly, if the 192.168. ∗ .∗ range were selected, 65 536 ($=2^{16}$) IPv6 addresses could be mapped. In our example, we select the reserved Class E network for our mapping since all three RFC 1918 address ranges were in use in the data set (because of the anonymization process). Using the Class E range has the added benefit of concentrating all IPv6 traffic at the end of the IPv4 address range, rather than embedding it between used ranges. It also offers 28 bits of unused IPv4 space. Finally, each indexed IPv6 address is translated (mapped) into its corresponding IPv4 address, depending on the range selected. In our example, the first and smallest IPv6 address becomes 240.0.0.0. The visualization is then generated using the existing IPv4-based tool, displaying both IPv4 addresses and mapped IPv6 addresses.

Figure 4 shows the resulting visualization displaying the traffic from the 15-minute data set. In Figure 4 (left), only IPv4 TCP and UDP connections are plotted. The integer source IP on the x-axis ranges from 0 to about 4.2 billion ($=2^{32}$), but there is a large section of unused addresses to the right of the graph. In Figure 4 (right), the IPv6

traffic is included as well, making use of the free space, thus displaying all TCP and UDP traffic from the data set of Section 'Dual-stack data set'. We can see in Figure 4 (right) that some trends, such as a constant stream of UDP connections to destination port 53 (DNS), are common to both the IPv4 and IPv6 portions of the graph. We also note that there is no HTTPS traffic (TCP port 443) over IPv6, which is common in IPv4, and practically no traffic to destination ports greater than 200 was recorded in the IPv6 address space.

This approach has several benefits. Current IPv4 visualization tools do not have to be modified to display dual stack captures after they are parsed using this technique. This allows forward compatibility and some breathing room while newer IPv6 visualization techniques are developed. Also, analysts can use a single visualization that displays all traffic, rather than doubling their workload by analyzing IPv4 and IPv6 data separately. When all traffic is on a single display, correlating activity of hosts using different IP versions is more straightforward.

As this approach moves one address space into another, analysts must be informed regarding this approach, or otherwise educated to differentiate between the two sections (the IPv4 range and the IPv6 range) of the graph. Failure to understand how the space is divided may otherwise lead to inaccurate interpretation of graphs. Finally, as IPv6 gains traction and IPv4 is phased out, our approach will need to be modified to support a larger number of IPv6 addresses. This is discussed in more detail in the following section.

Using IPv6 address hierarchy with treemaps

IPv6 address allocation standards dictate that addresses must be assigned in a hierarchical manner because of the
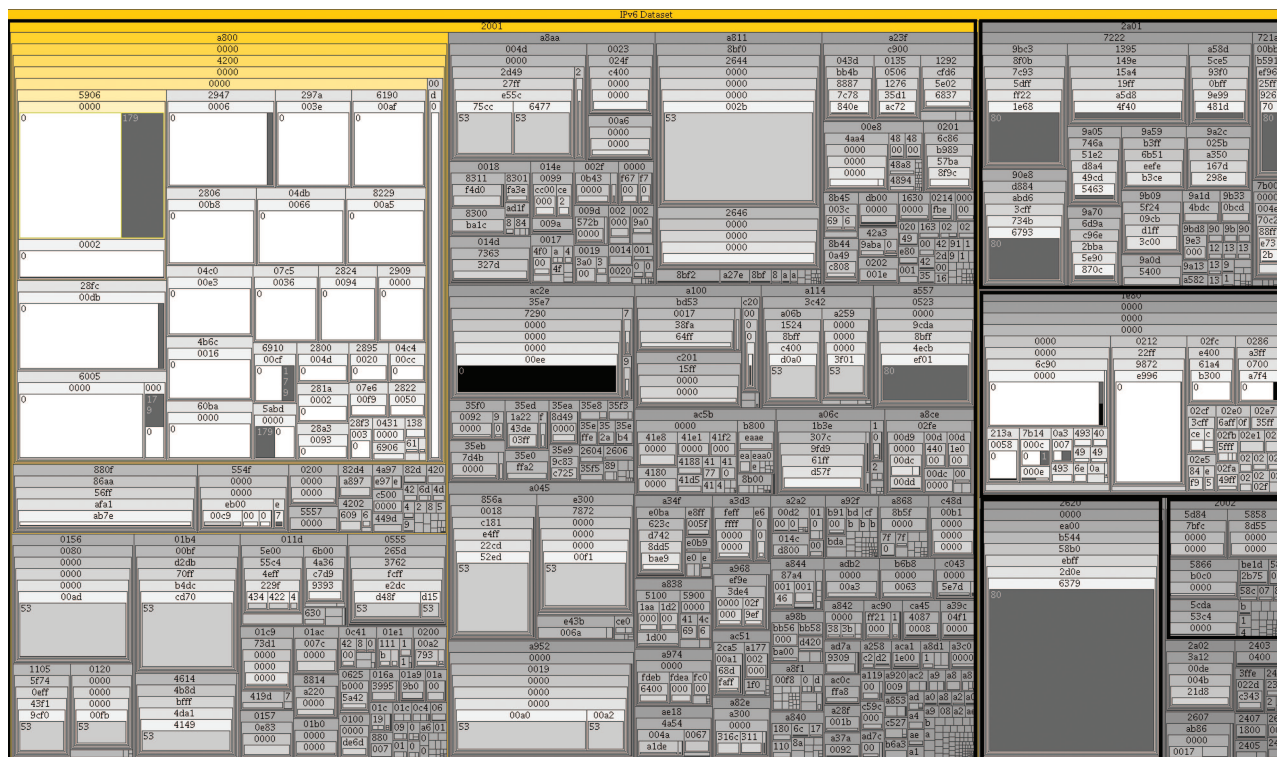
Figure 5: A treemap representing IPv6 source address, destination port and packet count (ICMP (white), TCP (dark grey), UDP (light grey), everything else (black)). Best viewed in color. Full resolution version available at: http://bit.ly/vizsec1.

extremely large address space. Assigning IPv6 addresses otherwise (for example, in an *ad hoc* way) would lead to huge routing tables on backbone Internet routers.

RFC 3587[18] specifies a global unicast format for aggregating IPv6 addresses to keep global routing tables efficient. In this format, bits of the address are grouped together, forming a three-level hierarchy: the global routing prefix, the subnet ID and the interface ID. Moving down the hierarchy (that is, reading more bits of the IPv6 address from left to right) reveals more about the IP address. The global routing prefix might tell us what regional registrar owns the address block, while the subnet ID might tell us what country and ISP own the sub-block. Finally, the interface identifier will uniquely identify a host or endpoint.

One way to visualize hierarchical data effectively is a treemap.[19] Treemaps have been used in the past to represent hard drive space (by using the file directory structure as hierarchy). An implementation of the treemap algorithm in Java was written by the HCI Laboratory at the University of Maryland.[20] The application reads in a text file in a specific format and outputs a treemap that can be navigated with dynamic tree height selection, node filtering and different layout types. We used this tool to illustrate how treemaps can be used to visualize IPv6 remote addresses.

Figure 5 shows an example treemap created from the data set of Section 'IPv6-only data set' (IPv6-only). The treemap represents all remote IPv6 addresses seen during the 4.5-hour capture period, the number of packets matching each port, and their protocol. The size of each rectangle is proportional to the number of packets, and the color represents the protocol. The treemap is hierarchical (using each hextet of the remote IPv6 address as a level), so viewing the first three hextets would display only the global routing prefix, and viewing the first four would display the subnet ID as well. The user can zoom in as far as the last hextet of the address, with the drawback of a more cluttered treemap. Rectangles that are too small to display a label can be expanded to full screen when clicked. Mousing over any rectangle gives a popup label displaying the packet count.

The treemap algorithm has generated five large rectangles (marked with black borders in Figure 5 corresponding to the first 16 bits of the IPv6 address–2001, 2A01, 2620 and 2A02–which are Internet Assigned Numbers Authority (IANA) assigned prefixes, and FE80 which is a link local block. The analyst can zoom in to the tree by double clicking on any of the labels, at which point the treemap will be redrawn using the selected node as the root of the tree. The tree is also searchable which might allow analysts to look for specific prefixes or network
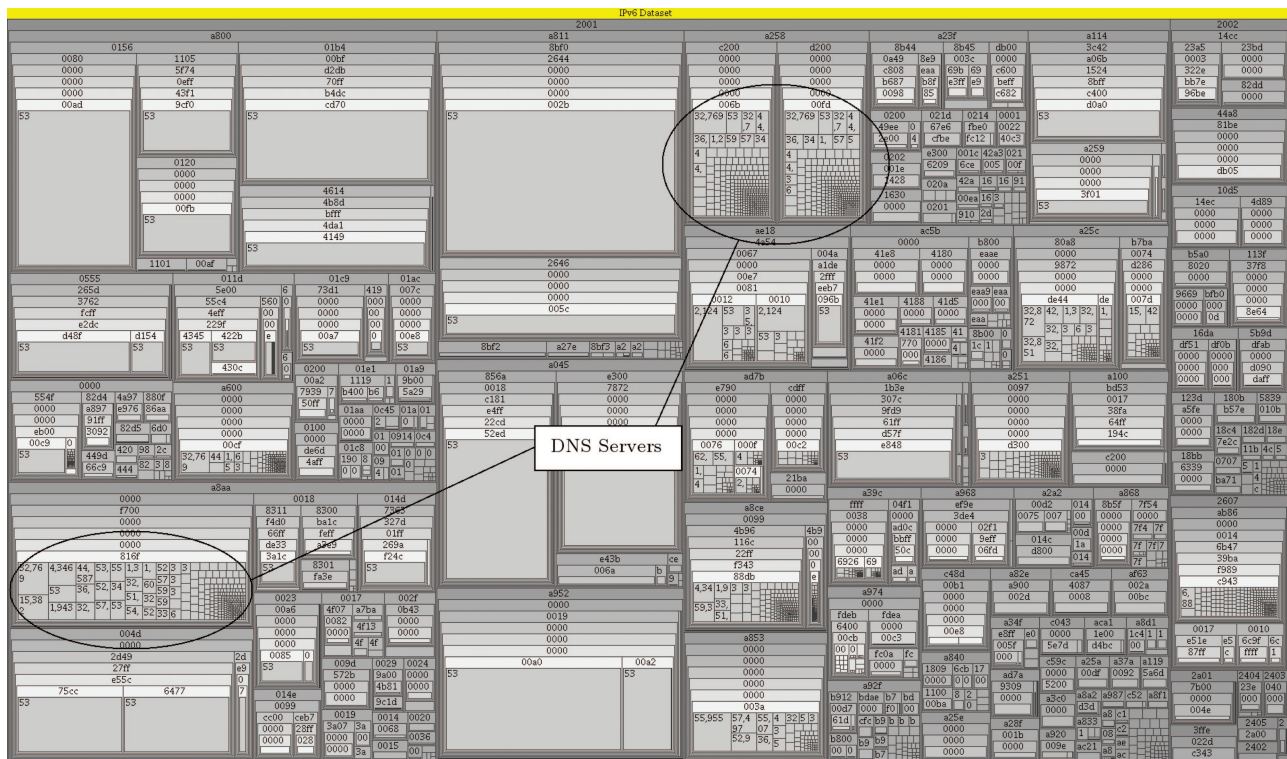
Figure 6: Finding DNS servers. Best viewed in color. Full resolution version available at: http://bit.ly/vizsec2.

blocks. Correlating this data with the list of IPv6 unicast assignments[21] by IANA reveals the geographical location of nodes.

One of the benefits of treemaps is that the user can quickly find sources behaving similarly by looking at the contents of nodes. For instance, in Figure 5 to the left, we notice that the highlighted hosts (mostly white nodes in the treemap) were seen sending mostly ICMPv6 traffic during the capture period. Some of these hosts were also seen sending TCP traffic on port 179, corresponding to the Border Gateway Protocol. These hosts are therefore probably routers or monitoring devices (or devices being monitored).

Another example of using treemaps for detecting specific types of hosts is presented in Figure 6. In this treemap, all hosts from our data set, except those seen generating UDP traffic, have been hidden. The resulting treemap displays three nodes with a distinctive destination port layout. All three nodes show that the most common destination port contacted was 32769 UDP (the second ephemeral port available by default on Unix-based operating systems) followed by 53 UDP and then many other high-numbered UDP ports. At a first glance, the contents of their nodes would suggest the three hosts behave similarly. This intuition is confirmed by searching for these IPv6 addresses using the SiLK tools, revealing that all three hosts only responded to requests on port 53 during the 4.5-hour capture.

Limitations

The three previously described techniques have advantages and drawbacks. The first method (removing unused addresses) will work well with data sets that have a small number of IPv6 hosts, and therefore has questionable scalability. The data sets of Sections 'IPv6-only data set' and 'Dual-stack data set' contained only 1924 and 3029 source IPv6 addresses, respectively, which generated uncluttered visualizations. In IPv4, we know that there can be at most 2^{32} remote addresses (actually far less due to reserved addresses) and that network address translation devices may hide a large number of hosts behind a single IP address. In IPv6, each Internet-enabled device can have its own address, potentially saturating visualizations created with this technique.

Real-time visualizations can be slow to implement with the first and second methods since they require a list to be sorted each time a new IP address is added. The treemap method of Section 'Using IPv6 address hierarchy with treemaps' is faster to update real-time, but requires the user to interact with the tool to reveal hosts with low volumes of traffic (for example, low and slow scans).

The transition mechanisms described in Section 'IPv6 review' will also require the analyst to further investigate points of interest, since source and destination IP addresses might not represent the end device, but rather a tunnel endpoint. Address ranges reserved for

transition mechanisms also appear to be more densely populated.

As IPv6 becomes more widely deployed, techniques like that described in Section 'Mapping IPv6 into unused IPv4 address space' will not scale well. While it is impossible to predict when these techniques might become obsolete, they can be further improved by allowing a many-to-one mapping of IPv6 to IPv4 addresses, and using additional complementary tools to display whether a point on the graph represents a single address or a set. Such methods will allow existing visualization techniques to continue to be used while the transition from IPv4 to IPv6 takes place.

Conclusion and Future Work

As IPv6 infrastructure is deployed, attackers are able to use IPv6 and have enabled support for it in their applications. Current visualization techniques fail to offer methods for displaying the source of connections if IPv6 is being used. This allows attackers the opportunity to remain undetected. We have proposed three new techniques that allow analysts using security visualization tools to see IPv6 network traffic, without requiring a major rewrite of the existing tools. While these methods seem to give good results on sample data sets, their practical utility once IPv6 deployment becomes mainstream remains unclear. A main objective has been to raise awareness of the need for security visualization tools to support IPv6 data sets. Future work includes finding new visualizations techniques for IPv6 traffic and improving the current proposed techniques.

References

1 Barrera, D. and van Oorschot, P.C. (2009) *Security Visualization Tools and IPv6 Addresses*. Proceedings of the 6th International Workshop on Visualization for Cyber Security (VizSec '09). Atlantic city, NJ, USA: IEEE, pp. 21–26.

2 Huston, G. (2009) IPv4 address report. http://www.potaroo.net/tools/ipv4/index.html, accessed 15 November 2009.

3 Van Riel, J.P. and Irwin, B. (2006) *InetVis, A Visual Tool for Network Telescope Traffic Analysis*. Proceedings of the 4th International Conference on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa (Afrigaph '06). New York, USA: ACM, pp. 85–89.

4 Lakkaraju, K., Yurcik, W. and Lee, A.J. (2004) *NVisionIP: Netflow Visualizations of System State for Security Situational Awareness*. In: Proceedings of the 2004 ACM Workshop on Visualization and Data Mining for Computer Security (VizSEC/DMSEC '04). New York, USA: ACM, pp. 65–72.

5 Koike, H., Ohno, K. and Koizumi, K. (2005) *Visualizing Cyber Attacks Using IP Matrix*. In: IEEE Workshop on Visualization for Computer Security (VizSec '05); pp. 91–98.

6 Deering, S. and Hinden, R. (1998) RFC 2460 – Internet Protocol, Version 6 (IPv6) Specification, http://www.faqs.org/rfcs/rfc2460.html, accessed 23 September 2009.

7 Gunderson, S. (2008) Measuring the current state of IPv6 for ordinary users, http://bit.ly/QTEXA, accessed 15 November 2010.

8 Moore, H. (2008) Exploiting tomorrow's internet today: Penetration testing with IPv6, October 2008, http://www.uninformed.org/?v=10&a=3, 23 September 2009.

9 IPv6 attack toolikit (2009) The hacker's choice http://freeworld.thc.org/thc-ipv6/, accessed 23 September 2009.

10 Taylor, T., Paterson, D., Glanfield, J., Gates, C., Brooks, S., and McHugh, J. (2009) *FloVis: Flow Visualization System*. Conference for Homeland Security: Cybersecurity Applications & Technology (CATCH '09). Washington DC: IEEE, pp. 186–198.

11 Nakamae, S., Sekiya, Y. and Murai, J. (1999) *A Study into a Visualization of an IPv6 Network*. In: Proceedings of the 9th Annual Conference of the Internet Society (INET'99). San Jose, CA.

12 Cooperative Association for Intimal Data Analysis (CAIDA) (2008) Visualizing IPv6 AS-level internet topology 2008. http://www.caida.org/research/topology/as_core_network/ipv6.xml, accessed 20 April 2009.

13 SiLK tools (2009) CERT network situational awareness team http://tools.netsa.cert.org/silk/, accessed 3 May 2009.

14 McHugh, J. (2009) *Data Structures for IPv6 Network Traffic Analysis Using Sets and Bags*. In: Proceedings of FloCon 2009, Scottsdale, AZ.

15 Measurement and Analysis on the WIDE Internet (2009) MAWI Working Group Traffic Archive. http://tracer.csl.sony.co.jp/mawi/, accessed 10 October 2009.

16 Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. and Lear, E. (1996) RFC 1918 – Address Allocation for Private Internets, http://www.faqs.org/rfcs/rfc1918.html, accessed 8 April 2009.

17 Kirkpatrick, S., Stahl, M. and Recker, M. (1990) RFC 1166 – Internet numbers, http://www.faqs.org/rfcs/rfc1166.html, accessed 8 April 2009.

18 Hinden, R., Deering, S. and Nordmark, E. (2003) RFC 3587 – IPv6 global unicast address format, http://www.faqs.org/rfcs/rfc3587.html, accessed 8 April 2009.

19 Shneiderman, B. and Wattenberg, M. (2001) *Ordered Treemap Layouts*. In: Proceedings of IEEE Symposium on Information Visualization (INFOVIS '01), pp. 73–78.

20 Treemap (2004) University of Maryland Human Computer Interaction Lab, http://www.cs.umd.edu/hcil/treemap/, accessed 23 September 2009.

21 IANA (2008) IANA IPv6 unicast address assignments, 13 May http://www.iana.org/assignments/ipv6-unicast-address-assignments, accessed 12 April 2009.