The ISP Column
A monthly column on things Internet

November 2013
Geoff Huston

IP Addresses and Traceback

This is an informal description the evolution of a particular area of network forensic activity, namely that of *traceback*. This activity typically involves using data recorded at one end of a network transaction, and using various logs and registration records to identify the other party to the transaction. Here we'll look at the impact that IPv4 address exhaustion and IPv6 transition has had on this activity, and also note, as we explore this space, the changing role of IP addresses within the IP protocol architecture.

Introduction

Public communications services are typically not completely anonymous facilities. Indeed, it is hard to comprehend how the public interest would be served if a public communications service operated in a totally opaque mode where every action was anonymous and untraceable, and use was unconstrained. Due to perhaps the darker side of human nature, such totally anonymous services tend to fall prey to abuse, overuse, leading to collapse of service integrity and a collapse of the service itself.

Therefore, it should come as no surprise that public communications systems typically include some aspect of end point identification and end point accounting of the use of the network. This record keeping may be limited to some form of summary of each user's use of the network, or may include so-called meta-data that includes information about the parties each user communicated with. It may even go as far as holding a copy of the communications itself, as has been the practice for some mobile operators and their retained records of SMS messages. In telephony the end point identification has been based on the device, not the user, and the device is identified by its telephone number. The accounting is typically based on the called number and the call record meta-data includes the time of day and call duration.

Given the Internet's role in public communications, how is a similar form of record keeping undertaken? How is traceback performed on the Internet? And how is this task changing as the internet itself evolves?

One way to look at this is to take a series of "snapshots" of the traceback operation, looking at how this achieved, and the assumptions that lie behind the traceback within the context of each "generation" of Internet service architecture.

The End-to-End IP Internet

We can start by looking at the Internet while it was still largely a research project, well before it assumed any role as a public communications utility, which would correspond to the late 1980's or thereabouts. In that architecture of the network the IP address of a connected device was its identification "key". If a ftp server logged the IP addresses of the clients who used its services, then this IP address would be sufficient to traceback to the end device. The essential information required to complete the traceback was the IP address registry data, that recorded the identity and contact details of the end user entity who received the address allocation.

An example of how this fits together is shown below.

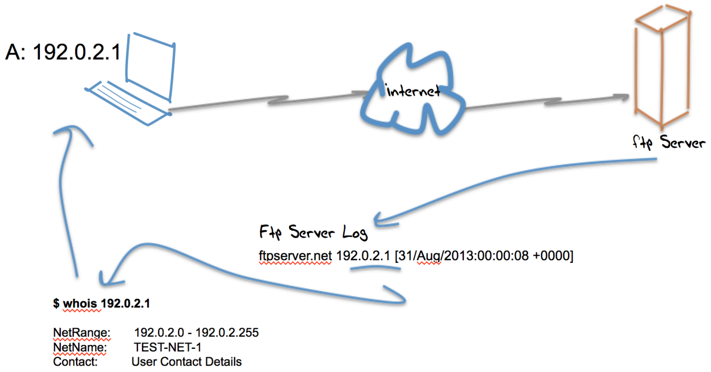

```
ftpserver.net 192.0.2.1 [31/Aug/2013:00:00:08 +0000]
```

```
$ whois 192.0.2.1

NetRange:     192.0.2.0 - 192.0.2.255
NetName:      TEST-NET-1
Contact:      User Contact Details
```

Figure 1 – Traceback in the End-to-End Internet

The server needs to record the remote IP address, and presumably the time of the access (although the time of access is not necessarily relevant to the traceback operation itself). The client's IP address can be used to perform a query against the address allocation registry, which, in turn, would provide the details of the end site to whom the IP address was allocated.

The basic attributes of the Internet at the time that allowed this simple form of traceback are:

- Each connected end site uses a stable public IP address range.
- Each address range is recorded in a registry, together with the end user contact information.
- Service providers route the addresses that the registry allocated to the end sites. Providers did not perform separate address allocation.
- Each end device was manually configured with a stable IP address.

There are strong similarities between this service model and that of the telephone network, and equally strong parallels between the underlying roles of the telephone number and the IP address in this context as endpoint identifiers.

The Introduction of Edge NATS

A major change to this model occurred with the introduction of Network Address Translators into edge firewalls and into various forms of customer premises network attachment devices. This was not the only change that occurred in the 1990's, but it was an indicator of a relatively fundamental change in the Internet's basic architecture.

The shift came in the form of a change in the service model of the Internet service provider sector, an evolution of a previous model of providing a basic packet transit service to capable end sites that already had obtained their own IP addresses, to a model of providing a connection service that included both the provision of addresses as well as packet transit services. The early forms of address scarcity pressures were also evident in this model, as the predominant addressing plan was to provision a single address to a customer connection, and do so as they connected to the access network, and then have the customer's equipment perform a NAT function to share this single address across all devices in the customer's network.

This introduces a few additional factors into the traceback task that did not exist in the previous model. The IP address is no longer directly associated with an end customer, but is registered against the service provider. An individual IP address may be used by multiple customers over time, so traceback to an individual customer requires using the logs generated by the service provider's authentication and accounting function (typically these would take for form of *Radius* logs) to determine which customer was using an address at a particular time. Also, an address is not necessarily bound to a single end device, but may be shared across multiple devices. So the address is now best seen as a ephemeral network identifier, and additional information is required to map visible use of the address to a particular end device.

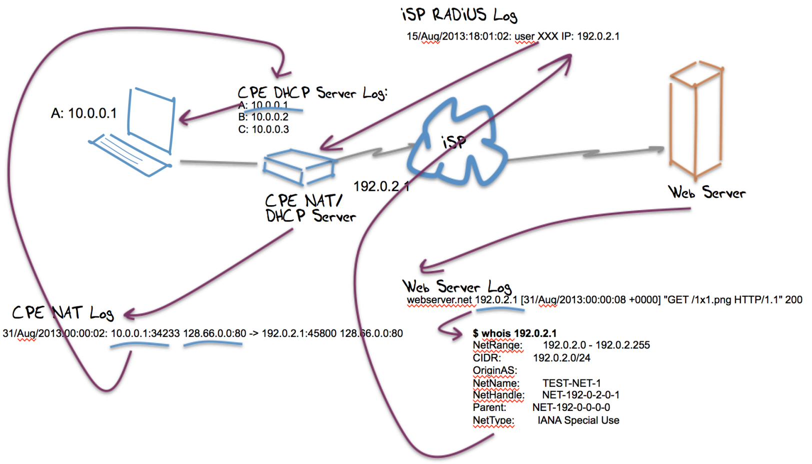

```
15/Aug/2013:18:01:02: user XXX IP: 192.0.2.1
```

```
A: 10.0.0.1
B: 10.0.0.2
C: 10.0.0.3
```

```
webserver.net 192.0.2.1 [31/Aug/2013:00:00:08 +0000] "GET /1x1.png HTTP/1.1" 200
```

```
31/Aug/2013:00:00:02: 10.0.0.1:34233 128.66.0.0:80 -> 192.0.2.1:45800 128.66.0.0:80
```

```
$ whois 192.0.2.1
NetRange:      192.0.2.0 - 192.0.2.255
CIDR:          192.0.2.0/24
OriginAS:
NetName:       TEST-NET-1
NetHandle:     NET-192-0-2-0-1
Parent:        NET-192-0-0-0-0
NetType:       IANA Special Use
```

Figure 2 – Traceback in the Edge-NAT Internet

Again, the server needs to record the remote IP address, and the time of the access (in this scenario the time is a necessary part of the traceback function). The client's IP address can be used to perform a query against the address allocation registry, which, in turn, would provide the details of the Internet Service Provider (ISP) to whom the IP address was allocated. Using the IP address and the time information the ISP's *Radius* logs should point to the corresponding customer account. But at this point the traceback operation typically stops. In theory, if the NAT device was recording and storing logs of its connection bindings, then by using the IP address of the server and the time of day it would conceivably be possible to match the exterior address to an interior address within the edge network, but it is uncommon to see such devices record comprehensive logs of their NAT address bindings.

None of the assumptions used in the previous model hold in this model of the Internet access infrastructure.

The new set of assumptions in this model are:

- The ISP operates an address pool
- Each end site is dynamically assigned a single IP address upon login (AAA)
- The edge site is dynamically addressed using a private address range and a DHCP server
- The single public address is shared by the private devices through a CPE NAT

Addresses are not registered against the edge network and not used individually used to identify host devices. In other words, addresses are no longer stable endpoint identifiers. The implication is that additional information is required to perform a reconstruction of the association of the address to an individual end device. The major additional information required to perform this traceback is the time of day, used in conjunction with the IP address as a lookup into the service provider's dynamic address configuration logs.

Mobility and IPv4 Address Exhaustion

This model of the Internet has been further refined by the rise of a new class of connections, namely the mobile Internet device.

The mobile service provider model has been built along slightly different lines to that of the wired network. Instead of a connectivity model that equates individual customers to edge networks, the mobile service model sees customers as individual devices, each of which is identified by a SIM card.

The scale of deployment of these devices, of the order of hundreds of millions of devices, rapidly overwhelmed the capacity of the IPv4 address plan, and the mobile data industry quickly adopted the use of NATs as an integral part of their deployment model. The critical difference between this model of NAT deployment and the previous model is that the NAT in this case is operated by the service provider, not by the customer. The generalized form of this approach to NATs has been termed "Carrier Grade NATs", (or CGNs) to distinguish them from the CPE-based edge NATs.

The other aspect of this rapid growth in the mobile service sector is that the IPv4 address plan has been effectively exhausted. The central pool of IPv4 addresses, operated by the Internet Assigned Names and Numbers Authority (IANA) was exhausted in February 2011. The Regional Internet Registry for the Asia Pacific was exhausted in April 2011, and the European and the Middle East Registry exhausted its pool of IPv4 addresses in September 2012. The Regional Registries for North and South America are expected to exhaust their inventories of IPv4 addresses in the coming months.

The long term plan to respond to this exhaustion of IPv4 addresses was to migrate the entire Internet to a new protocol, namely IPv6. However IPv6 is not backward compatible with IPv4, which means that for a protracted period the Internet will be operating using two protocols simultaneously. But having to continue to stretch a depleted IPv4 address pool across an ever-expanding Internet is now forcing service providers to look at methods that stretch the IPv4 address resources over ever-larger spans of customer networks. One approach is to take the Carrier Grade NATs of the mobile service provider sector and use them in the wired networks. This is viable because there is no technical impediment for connections to traverse multiple NATs in a single path.

The traceback function in this form of network is more complex, as shown in the figure below.

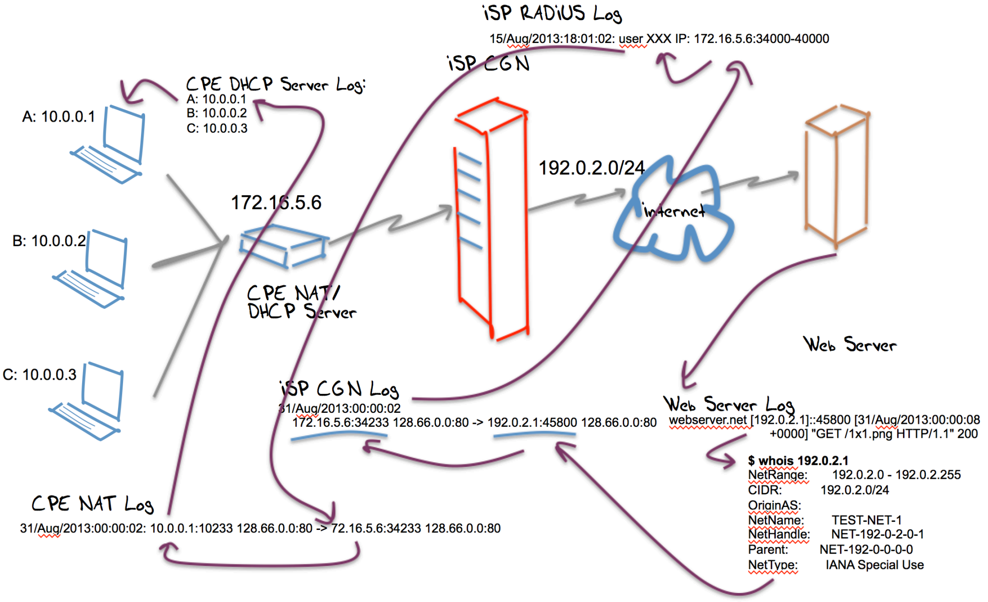

```
15/Aug/2013:18:01:02: user XXX IP: 172.16.5.6:34000-40000
```

```
31/Aug/2013:00:00:02
 172.16.5.6:34233 128.66.0.0:80 -> 192.0.2.1:45800 128.66.0.0:80
```

```
webserver.net [192.0.2.1]::45800 [31/Aug/2013:00:00:08
          +0000] "GET /1x1.png HTTP/1.1" 200
```

```
$ whois 192.0.2.1
NetRange:      192.0.2.0 - 192.0.2.255
CIDR:        192.0.2.0/24
OriginAS:
NetName:       TEST-NET-1
NetHandle:     NET-192-0-2-0-1
Parent:      NET-192-0-0-0-0
NetType:      IANA Special Use
```

```
31/Aug/2013:00:00:02: 10.0.0.1:10233 128.66.0.0:80 -> 72.16.5.6:34233 128.66.0.0:80
```

Figure 3 – Traceback in the CGN + Edge-NAT Internet

The first major change here is that the server logs now need to record the remote client's IP address and their transport protocol port number, as well as the time. When CGNs are used it is now necessary to have the complete connection profile (source and destination IP address and port numbers) as well as a time value that is accurate to the level of milliseconds. It is noted that few default service configurations include the recording of the remote port used in service connections, which is a critical weakness in this scenario, and keeping accurate time is also not as widespread across services as it could be.

At this point the client side IP address can be used as a registry lookup key, and this will point to the provider. In the case where there is a CGN in the path, the next step in the traceback process is to use the connection profile of source and destination addresses and port numbers and the time value as a lookup into the carrier provider's logs of the CGN bindings, and find the equivalent local IP address and port number that the CGN had used at the time for this connection. This local addresses is then used to look up the service provider's Radius logs to find the customer who was assigned this local-side addresses, which then points to the customer.

Of note here is that IP addresses have lost all their association as a direct reference to an individual device or even to an individual customer. They have become ephemeral connection tokens that resolve to the level of a service provider but no further.

It is also worth noting that the logs generated by a CGN can be quite large, and their capture, storage and lookup can represent a non-trivial challenge. Cable Labs in the US has reported (Chris Grundeman, http://www.nanog.org/meetings/nanog54/presentations/Tuesday/GrundemannLT.pdf) on a measurement exercise in scoping the size of these CGN logs, and using an approach that required between 150 – 450 bytes of log information per individual CGN binding operation, and a volume of between 33,000 to 216,000 such binding operations per customer per day, which represents a data archival log load of some 5 – 96 Mb per day. For a SP with 1M customers, this is potentially 1Pb of log data per month.

It has been argued that this example is an extreme case, and various forms of user port blocks and similar methods of port grouping can reduce the logging load arising from CGN use (as can data compression applied to the log files), but at the same time it must be considered that the higher the level of address compression the fewer the opportunities that exist for various forms of port blocking and aggregation. For example, if one were to use a CGN that assigned 16 customers per public IP address, then each customer would have 4,096 ports for use in each of TCP and UDP, rather than the 65,536 ports that is used with the CPE-based NAT architecture. Increasing the number of customers per public address decreases the number of ports, which in turn impacts on the integrity of applications. If one were to eliminate the entire concept of port reservations and moved to an open port contention model then up to 1,000 customers could be provisioned against each public IP address, but at that point the logging overhead reverts back to the high volumes referred to previously. Yet higher address utilization densities can be achieved using an approach termed "5-tuple NATTing", where the NAT can use the same public IP address and TCP or UDP port address in multiple connections simultaneously, as long as the connections using these common addresses and ports are directed to distinct destination addresses or ports.

At this point the concept of an IP address loses all semantic significance other than as an identification of the provider. Its residual role is as an element of an ephemeral session identifier, and little else. The session identification key is becomes the complete IP connection set, which is composed of the source and destination addresses, the transport protocol identifier and the source and destination port addresses, as used in the connection.

Dual Stack and IPv6 Transition

Unfortunately the potential complexities in this picture do not stop at that point. With many providers seeing an inevitable conclusion of this transition being in the deployment of an IPv6 service platform, some thought has gone into ways in which an IPv4 service can be provided to customers across an access network that has been already concerted to an IPv6-only service platform. The DS-Lite approach (RFC6333) is one such example, where the CGN performs the conventional local / public address and port binding, but in this case the resultant IPv4 packet is then encapsulated into an IPv6 packet header where the IPv6 addresses identify the CGN and the customer CPE units.

This can be further extended by replacing the encapsulation by a direct translation function, where theIPv6 addresses can be used to both identify the CGN and the CPE, but also hold the CGN-mapped IPv4 addresses, as is specified in 464XLAT (RFC6877).

Generalizing from these individual approaches to dual stack transition, it can be observed that there are a wide variety to protocol translation and encapsulation mechanisms that have been developed to assist in the transition of the Internet to IPv6.

One possible taxonomy of such transition technologies in shown in the following figure (Randy Bush, http://meetings.apnic.net/__data/assets/pdf_file/0016/45241/120229.apops-v4-life-extension.pdf)

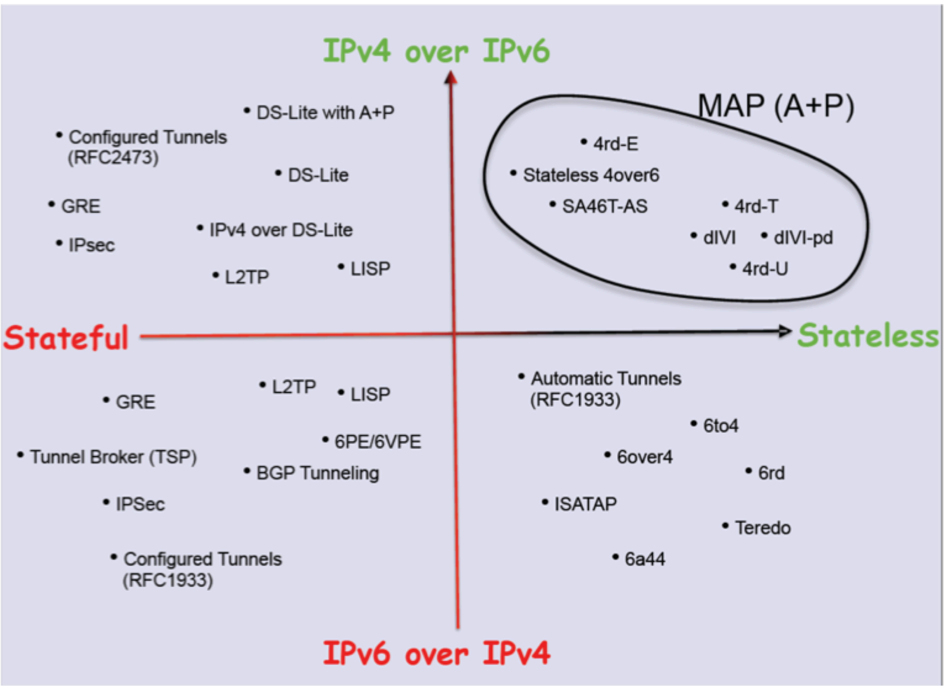

Figure 4 – Dual Stack Transition Technologies (R. Bush)

It is becoming clear that in this Internet, an Internet that now uses CGNs operating in various modes, and an Internet that variously uses IPv6 transition technologies in diverse ways, is an Internet that no longer has a consistent interpretation of the role of addresses within its use. There is no longer a single service architecture in the Internet, as we see different providers respond in unique ways to differing commercial pressures and opportunities, which, in turn, sees these providers utilize differing technology in their networks. It also seems reasonable to observe that the longer we sit within this world of an depleted IPv4 address pool and at the same time have to field a transitioning dual stack environment, the greater the level of technical diversity and complexity that will be used in these access and service networks.

It appears that if one wishes to map an external IP address and time of day back to an end point then the record keeping requirements on carriage providers and service providers becomes more encompassing: for every active middleware element it would be necessary to hold the exact time and the exact set of address transforms that were applied to each IP connection flow. The questions this raises are whether such record keeping practices are affordable by the industry, given that the shift in granularity of the records is now shifting from a connection session to every individual TCP or UDP stream. And in an environment of continuing growth and intensity of use, it is also reasonable to ask whether such record keeping practices would be scalable into the future.

The traceback toolkit to perform any form of forensic investigation into past activity in the network would need to include:

- Precise time, source and dest IP addresses, protocol and port information
- Access to all ISP middleware logs
- Content Distribution Network Service Provider logs
- Network and Middleware deployment maps
- The V6 Transition technology map used by the ISP
- A thorough understanding of vendor's equipment behaviour for various applications
- A thorough understanding of application behaviours

A similar question exists here in questioning whether such resources are available for the folk who are tasked with performing such forensic exercises. At what point in time does the escalating diversity and complexity of these networks scale beyond any reasonable set of forensic capabilities?

An IPv6 Internet

The underlying long term plan for the Internet was to leave the IPv4 Internet behind, and transition the network to one that operates on IPv6. Obviously, the plan included the consideration that we were to achieve this before we exhausted the pools of remaining IPv4 addresses, but in any case its reasonable to ask what this forensic task looks like once we complete this transition. What does traceback look like in an all-IPv6 network?

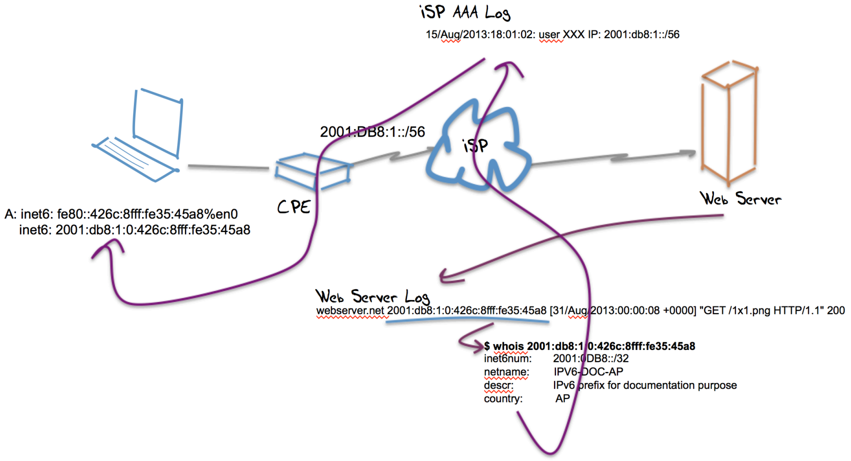

```
15/Aug/2013:18:01:02: user XXX IP: 2001:db8:1::/56
```

```
A: inet6: fe80::426c:8fff:fe35:45a8%en0
   inet6: 2001:db8:1:0:426c:8fff:fe35:45a8
```

```
webserver.net 2001:db8:1:0:426c:8fff:fe35:45a8 [31/Aug/2013:00:00:08 +0000] "GET /1x1.png HTTP/1.1" 200
```

```
$ whois 2001:db8:1:0:426c:8fff:fe35:45a8
inet6num:     2001:0DB8::/32
netname:      IPV6-DOC-AP
descr:        IPv6 prefix for documentation purpose
country:      AP
```

Figure 5 – Traceback in the IPv6 Internet

The original intention in IPv6 was to revert to a simpler network architecture that is more closely aligned to that of the original IPv4 Internet, where individual IPv6 addresses are uniquely associated with an end device. The addressing plan in IPv6 is still aligned to service providers so the registry entry for the address will still point to the service provider, and the service provider's logs will be needed to map the address to a site prefix that has been assigned to an end customer.

Some forms of local site auto-configuration have included use of the 48-bit MAC address in the low 64 bits of the IPv6 address, so it is feasible in this model to track back all the way to the end device.

However, IPv6 is also evolving. A number of widely deployed IPv6 implementations now use "privacy address", which turn the low 64 bits into a random bit sequence which is periodically refreshed. In this context there is no effective form of traceback to the device. The closest this exercise will get to the end point is the site prefix being used by the end site.

Even with privacy addresses, the IPv6 network paints a far more coherent picture of address use that the byzantine complexities that are emerging in today's Internet. With IPv6 privacy addressing in use it may not be necessarily possible to resolve a particular address back to a unique end point and a particular device, but such an exercise in an IPv6 Internet would be able to provide consistent visibility to the level of granularity of the end site prefix without the massive log generation, maintenance and data analysis factors that are associated with the use of CGNs and various dual stack transition mechanisms.

Privacy and Accountability

There is always a balance to be struck between the appropriate levels of privacy and accountability in any form of a public communications system.

A system that admits no form of privacy whatsoever rapidly falls into disuse when an alternative emerges that does offer basic protections to its users regarding the privacy of their communications. Additionally, there is much to be said in favour of the use of robust services that provide basic protection of the content of communications, including the use of adequate cryptography and careful management of cryptographic keys and judicious choice in the selection of trust anchors.

But that does not mean that an entirely opaque network where privacy is absolute in all aspects in necessarily in the public interest either. When assured anonymity is not only possible, but is the default mode of communication, it appears that we see the emergence of abuse and fraud. The level of toxic activity on today's Internet, with its overwhelming levels of abuse of mail, port scanning, vulnerability exploitation, bot armies, DOS attacks, phishing and the like, all point to a vibrant level of exploitation of the basic ability to perform such acts without ready attribution.

Where we are right now, with an Internet that increasingly uses ever more complex forms of transforms on IP packet headers in flight, and one that has managed to transform end point identifiers in the form of IP addresses into ephemeral session tokens is one that does not readily admit the basic levels of accountability.

Perhaps this is not a good place to be.

Disclaimer

The views expressed are the authors' and not those of APNIC, unless APNIC is specifically identified as the author of the communication. APNIC will not be legally responsible in contract, tort or otherwise for any statement made in this publication.

About the Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

www.potaroo.net