The ISP Column
A monthly column on things Internet

May 2013
Geoff Huston

A Royal Opinion on Carrier Grade NATs

There are still a number of countries who have Queen Elizabeth as their titular head of state. My country, Australia, is one of those countries. It's difficult to understand what exactly her role is these days in the context of Australian governmental matters, and I suspect even in the United Kingdom many folk share my constitutional uncertainty. Nevertheless, it's all great theatre and rich pageantry, with great press coverage thrown in as well. In the United Kingdom every year the Queen reads a speech prepared by the government of the day, which details the legislative measures that are being proposed by the government for the coming year. Earlier this month the Queen's speech included the following statement in her speech:

> "In relation to the problem of matching Internet Protocol addresses, my government will bring forward proposals to enable the protection of the public and the investigation of crime in Cyberspace."
> http://www.youtube.com/watch?v=UWwK3z3GvzY&amp;feature=youtube_gdata (5:45)

As the Guardian pointed out:

> The text of the Queen's speech gives the go-ahead to legislation, if needed, to deal with the limited technical problem of there being many more devices including phones and tablets in use than the number of internet protocol (IP) addresses that allow the police to identify who sent an email or made a Skype call at a given time.
> http://www.guardian.co.uk/politics/2013/may/08/queens-speech-snoopers-charter

What's the problem here?

The perspective of various law enforcement agencies is that the Internet is seen as a space that has been systematically abused, and too many folk are felling prey to various forms of deceit and fraud. If you add to that the undercurrent of concern that the Internet contains a wide range of vulnerabilities from the perspective of what we could generally term "cybersecurity," then it's not surprising to see law enforcement agencies now turning to legislation to assist them in undertaking their role. And part of their desired toolset in undertaking investigations and gathering intelligence is access to records from the public communications networks of exactly who is talking to whom. Such measures are used in many countries, falling under the generic title of "data retention."

In the world of telephony the term "data retention" was used to refer to the capture and storage of call detail records. Such records typically contain the telephone numbers used, time and duration of the call, and may also include ancillary information including location and subscriber details. Obviously such detailed use data is highly susceptible to data mining, and such call records can be used to identify an individual's associates and can be readily used to identify members of a group. Obviously, such data has been of enormous interest to various forms of law enforcement and security agencies over the years, even without the call conversation logs from direct wire tapping of targeted individuals. The regulatory measures designed to protect access to these records vary from country to country, but access is typically made available to agencies on the grounds of national security, law enforcement or even enforcement of taxation conformance.

So if that's what happens in telephony, what happens on the Internet?

Here the story is a continually evolving one, and these days the issues of IPv4 address exhaustion and IPv6 are starting to be very important topics in this area. To see why it is probably worth a looking at how this used to happen and what technical changes have prompted changes to the requirements related to data retention for Internet Service Providers (ISPs).

The original model of the analogous data records for the Internet was the registry of allocated addresses maintained by Internet Network Information Centre, or Internic. This registry did not record any form of packet activity, but was the reference data that shows which entity had been assigned which IP address. So if you wanted to know what entity was using a particular IP address, then you could use a very simple “whois” query tool to interrogate this database:

```
$ whois -h whois.apnic.net 202.12.29.211
% [whois.apnic.net node-4]
% Whois data copyright terms    http://www.apnic.net/db/dbcopyright.html

inetnum:        202.12.28.0 - 202.12.29.255
netname:        APNIC-AP
descr:          Asia Pacific Network Information Centre
descr:          Regional Internet Registry for the Asia-Pacific Region
descr:          6 Cordelia Street
descr:          PO Box 3646
descr:          South Brisbane, QLD 4101
descr:          Australia
```

However, this model of the registry making direct allocations to end user entities stopped in the early 1990’s with the advent of the ISP. The early models of ISP service were commonly based on the dial-up model, where a customer would be assigned an IP address for the duration of their call, and the IP address would return to the free pool for subsequent reassignment at the end of the call. The new registry model was that the identity of the service provider was described in the public address registry, and the assignment of individual addresses to each of their dial-up customers was information that was private to the service provider. Now if you wanted to know what entity was using a particular IP address you also had to know the time of day as well, and while a “whois” query could point you in the direction of whom to ask, you now had to ask the ISP for access to their Access, Authentication and Accounting (AAA) records, typically the radius log entries, in order to establish who was using a particular IP address at a given time. Invariably, this provider data is private data, and agencies wanting access to this data had to obtain appropriate authorization or warrants under the prevailing regulatory regime.

This model of traceback has been blurred by the deployment of edge NATs, where a single external IP address is shared across multiple local systems serviced by the NAT. This exercise can therefore trace back to the NAT device, but no further. So with access to this data you can get to understand the interactions on the network at a level of granularity of customer end points, but not at a level of individual devices or users.

We’ve used this model of Internet address tracking across the wave of cable and DSL deployments. The end customer presents their credentials to the service provider, and is provided with an IPv4 address as part of the session initiation sequence. The time of this transaction, the identity of the customer and the IP address is logged, and when the session is terminated the address is pulled back into the address pool and the release of the address is logged. The implication is that as long as the traceback can start with a query that includes an IP address and a time of day, its highly likely that the end user can be identified from this information.

But, as the Guardian’s commentary points out, this is all changing again. IPv4 address exhaustion is prompting some of the large retail service providers to enter the Carrier Grade NAT space, and join what has already become a well established practice in the mobile data service world. The same week of the Queen's speech, BT announced a trial of Carrier Grade NAT use in its basic IP service (http://www.techweekeurope.co.uk/news/bt-retail-trials-ip-address-sharing-carrier-grade-na-115411).

At the heart of the Carrier Grade NAT approach is the concept of sharing a public IP address across multiple customers at the same time. An inevitable casualty of this approach is the concept of traceback in the internet and the associated matter of record keeping rules. It is no longer adequate to front up with an IP address and a time of day. That is just not enough information to uniquely distinguish one customer's use of the network from another's. But what is required is now going to be dependant on the particular NAT technology that is being used by the ISP. If the CGN is a simple port-multiplexing NAT then you need the external IP address and the port number. When combined with the CGN-generated records of NAT's bindings of internal to external address, this can map you back to the internal customer's IP address, and using the ISP's address allocations records, this will lead to identification of the customer.

So traceback is still possible in this context. In a story titled "Individuals can be identified despite IP address sharing, BT says" the newsletter out-law.com (produced by the law firm Pinsent Masons) reports:

> BT told Out-Law.com that its CGNAT technology would not prevent the correct perpetrators of illegal online activity from being identified.
>
> "The technology does still allow individual customers to be identified if they are sharing the same IP address, as long as the port the customer is using is also known," a BT spokesperson said in a statement. "Although the IP address is shared, the combination of IP address and port will always be unique and as such these two pieces of information, along with the time of the activity can uniquely identify traffic back to a broadband line. […] If we subsequently receive a request to identify someone who is using IP address x, and port number y, and time z we can then determine who this is from the logs," the spokesperson said. […] "If only the IP address and timestamp are provided for a CGNAT customer then we are unable to identify the activity back to a broadband line," they added.
> http://www.out-law.com/en/articles/2013/may/individuals-can-be-identified-despite-ip-address-sharing-bt-says/

But port-multiplexing NATs are still relatively inefficient in terms of address utilization. A more efficient form of NAT multiplexing uses the complete 5-tuple of the connection signature, so that the NAT's binding table uses a lookup key of the protocol field and the source and destination addresses and port values. This allows the NAT to achieve far higher address sharing ratios, allowing a single external IP address to be shared across a pool of up to thousands of customers.

So what data needs to be collected by the ISP to allow for traceback in this sort of CGN environment? In this case the ISP needs to collect the complete 5-tuple of the external view of the connection, plus the start and stop times at a level of granularity to the millisecond or finer, together with the end-user identification codes. Such a session state log entry takes typically around 512 bytes as a stored data unit.

How many individual CGN bindings, or session states, does each user generate? One report I've seen points to an average of some 33,000 connections per end customer each day. If that's the case then the implication is that each customer will generate some 17Mbytes of log information every day. For a very large service provider, with, say, some 25 million customers, that equates to a daily log file of 425Tbytes. If these CGN records were produced at an unrealistically uniform rate per day, that's a constant log data flow of some 40Gbps. At a more realistic estimate of the busy period peaking at 10 times the average, the peak log data flow rate is some 400Gbps.

That's the daily load, but what about longer term data retention storage demands? The critical questions here is the prevailing data retention period. In some regimes it's 2 years, while in other regimes it's up to 7 years. Continuing with our example, holding this volume of data for 7 years of data will consume 1,085,875 Terrabytes, or 1.0 Exabytes to use the language of excessively large numbers. And that's even before you contemplate backup copies of the data! And yes, that's before you contemplate an Internet that becomes even more pervasive and therefore of course even larger and used more intensively in the coming years.

The questions such a data set can answer also requires a very precisely defined question. It's no longer an option to ask "who used this IP address on this date?" Or even "who used this IP address and this port address in this hour?" A traceback that can penetrate the CGN-generated address overuse fog requires the question to include both the source and destination IP addresses and port numbers, the transport protocol, and the precise time of day, measured in milliseconds. This last requirement, of precise coordinated time records, is a new addition to the problem, as traceback now requires that the incident being tracked be identified in time according to a highly accurate time source running in a known timezone, so that a precise match can be found in the ISP's data logs. It's unclear what it will cost to collect and maintain such massive data sets, but its by no means a low cost incidental activity for any ISP.

No wonder the UK is now contemplating legislation to enforce such record keeping requirements in the light of the forthcoming CGN deployments in large scale service provider networks in that part of the world. Without such a regulatory impost its unlikely that any service provider would, of their own volition, embark on such a massive data collection and long term storage exercise. One comment I've heard is that in some regimes it may well be cheaper not to collect this information and opt to pay the statutory fine instead – it could well be cheaper!

This is starting to look messy. The impact of CGNs on an already massive system is serious, in that it alters the granularity of rudimentary data logging from the level of a connection to the Internet to the need to log each and every individual component conversation that every consumer has. Not only is it every service you use and every site you visit, but its even at the level of every image, every ad you download, everything. Because when we start sharing addresses we now can only distinguish one customer from another at the level of these individual basic transactions. Its starting to look complicated and certainly very messy.

But, in theory in any case, we don't necessarily have to be in such a difficult place for the next decade and beyond.

The hopeful message is that if we ever complete the transitional leap over to an all-IPv6 Internet the data retention capability reverts back to a far simpler model that bears a strong similarity to the very first model of IP address registration. The lack of scarcity pressure in IPv6 addresses allows the ISP to statically assign a unique site prefix to each and every customer, so that the service providers data records can revert to a simple listing of customer identities and the assigned IPv6 prefix. In such an environment the cyber-intelligence community would find that their role could be undertaken with a lot less complexity, and the ISPs may well find that regulatory compliance, in this aspect at least, would be a lot easier and a whole lot cheaper!

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.

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