

Geoff Huston  
April 2017

## ARIN 39 Report

Having just spent two and a half days at an ARIN Public Policy Meeting, I'd like to share some of my impressions of the meeting, and the state of address policy in the region served by ARIN. It's clear that ARIN has largely undertaken its mission with careful and considered distinction. It has achieved much of what is efficiently achievable in a consensus-driven open process, and left what is overly contentious, or too enmeshed in the vagaries of history to even attempt to unravel. What this means is that these days ARIN meetings are generally not filled with high drama and fraught contentious debate, and this latest meeting was no exception. In terms of policy, the meeting largely dealt with matters that were obvious, or could be seen as fine level clarification of existing policies.

However, there was one point where I thought that there was some level of disconnect, and that concerns the contents of the ARIN number registry. Over the years ARIN and its predecessors, have made some 50,000 allocations of IP addresses and/or Autonomous System Numbers, and these entities, and the resources they currently hold, are listed in the ARIN registry. However, that's not all that you can find in this registry. For many years, it has been a requirement to submit to ARIN the details of certain address sub-assignments made by these recipients of ARIN addresses. ARIN's database holds more than 700,000 records that record addresses and points of contact (PoCs). It is unclear to what extent this secondary set of records of sub-assignments is complete, or even whether it was accurate at any point in time. Such records were evidently reviewed at those times when the original address holder requested additional addresses and revised their sub-assignment records as part of the ARIN review of address utilization, but at other times the level of attention paid to the completeness and currency of these sub-assignment records was somewhat variable.

It was evident in the policy discussion at this meeting that a number of agencies, generally associated with aspects of law enforcement (LEAs), would like to see more attention paid to this registry, and the policy proposal essentially proposed to task ARIN with some level of continual activity that would monitor the accuracy of all of these records in the database.

It's certainly the case that almost all forms of criminal and extreme behaviours are "cyber crimes" of one form or another, and relating traces of online criminal or extremist behaviour to the identities of individuals is a natural desire by these LEAs. The analogous reference in the days of telephony was a reverse "white pages" where a phone number could be traced to a subscriber. If this was possible for the telephone network, why can't we do this for the Internet? After all everyone who generates an online transaction uses an IP address. Why can't we publish some form of "Internet Reverse White Pages" listing that associates all IP addresses with end entities? It was evident from the discussion at ARIN that some LEAs are interested to see this happen, and in the case of North America, they are keen to see ARIN take a leading role in facilitating this. And this registry of some 700,000 IP address holders is thought to be a really good place to start. This policy proposal is advocating that ARIN regularly audit all these PoCs and ensure that they are accurate.

That is all well and good, but there are some additional aspects to consider here. In the US alone there are estimated to be some 282 million individual users of the Internet. Which subset of this rather significant set of users is listed in the ARIN database as a PoC for an IP address? Obviously, there are a

massive number of assigned IP addresses where no PoC exists in ARIN's registry. These sub-assignment records are not records that were created by ARIN, and are not curated by ARIN. Indeed, it is probably the case that for many of these listed entities they are unaware that they are listed in this database. It does seem a little far-fetched to compel ARIN to contact a set of folk that have no relationship and potentially no knowledge of ARIN, and start a conversation about the accuracy of the information that ARIN holds that describes them and their contact details.

The obvious weakness of this database in terms of its level of comprehensive coverage of attribution of effective end user assignment of IP addresses is probably as much to do with the architecture of today's Internet as it is to do with any failings in attempting to keep these sub-assignment records up to date. Many retail access providers use either dynamic address assignment pools in those cases where public IP addresses are assigned to end users, or, more commonly these days, the use of public addresses is completely automated by virtue of carrier grade NAT deployments. In the latter case, not only is the address dynamically assigned, but it is likely to be shared over many customers, possibly up to tens of thousands of customers in a large CGN pool.

The overall numbers of today's Internet graphically illustrate the scope of address sharing. Common current estimates indicate that the Internet is populated by between 12 to 14 attached devices, and at the same time we estimate that between 1 1/2 and 2 billion IP addresses are used by these devices. Obviously not every endpoint has its own unique IP address. Maybe its time to walk away from phone books, and walk away from the concept that there is some underlying persistence in the association of individual IP addresses and connected end point devices.

It's certainly reasonable for a registry such as ARIN, or any of the other four Regional Internet Registries, to work diligently to ensure that the data in their registry that relates to address assignments directly made by the registry is complete and accurate at all times. But it is perhaps not so reasonable to compel these same registries to create a public repository of sub-assignment of addresses and the related record keeping of dynamic address assignments by service providers. The registry has little in the way of effective inducement or enforcement ability to ensure that any such records are complete, current or accurate, and partial data sets of dubious provenance are often of less value than having no data in the first place. What might help here is for ARIN might be to very clearly mark all data that relates to address assignments made by ARIN, and ensure that such data is actively curated by ARI, even to the extent of being able to query the ARIN-only entries for address records. As for the other 700,000 or so entries, maybe the case can be made that no data at all is better than incomplete bad data!

In another policy discussion the long standing debate over address transfer policies was raised. In ARIN there continues to be a school of thought that believes, strongly, that a recipient of an address transfer needs to be able to meet some "demonstrated need" criteria before the transfer will be recorded in ARIN's registry. There is also a school of thought that believes, strongly, that the imposition of policies that prevent the registration of address transfers does not prevent the transfer, but instead disconnects the registry from the "ground truth" of the network itself, demeaning the utility of the registry as a common reference source relating to the current disposition of addresses. There is no commonly acceptable resolution to this debate that has emerged so far, and certainly not at ARIN 39! Instead, we are seeing some tinkering of the very fine level details of aspect of handling address transfers in a couple of the policy proposals discussed at this meeting. No doubt this topic will be revisited at future ARIN meetings.

As is usual for ARIN, ARIN 39 was a well organised meeting, fulfilling ARIN's undertaking to support an open and transparent policy development process. The meeting was well organised, well supported for both local and remote participants and the efforts to ensure that all participants were well briefed on the matters under consideration were nothing short of exemplary. For this both ARIN, and the participants at these public address policy meetings deserve plaudits in undertaking an important and at times difficult task with friendliness and a common desire to seek a working consensus wherever and whenever that's achievable.

---

## Author

*Geoff Huston* B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

*[www.potaroo.net](http://www.potaroo.net)*

---

## Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.