Hunting the Bogon

May 2004

There were a number of "adventure" games in the late 1970's and early 80's. By today's standards the original versions of these computer games were not just primitive, they make banging the rooks together look sophisticated. But at the time they were a pointer to the use of computers as something more than just numerical calculators, or business machines. Programs could be imaginative and even fun (well ok, 'fun' was different then!). One of the more widespread early versions of this kind of game was "Hunt the Wumpus'. The web being what it is these days you can meander back to 1976 at

http://www.atariarchives.org/bcc1/showpage.php?page=247

For those who want a first-hand experience of 1970's computer games, check out: http://www.taylor.org/~patrick/wumpus/ As far as I can tell the only difference is that the original games proceeded at a pace that could only be described as glacial!

What's a "Bogon"?

Good question. The word *bogon* probably has lots of meanings in various contexts, but in the context of the Internet address realm a *bogon* refers to the use of an address or, more generally a route object, that is not duly authorized by the entity to which the address, or resource, was originally assigned. I must admit that I've yet to hear of any other word that succinctly takes that rather long- winded phrase and sums it up with one word!

There are two kinds of bogon objects in the inter-domain space - the first is the advertisement of IP addresses, and the second is the use of Autonomous System numbers within the AS Path attribute.

The problem that bogons present is generally related to threats to the integrity of the Internet's address space. Bogons are what could be considered to be 'unauthorized' use of the address space. In one form of bogon there is no record of the original resource allocation ever having been made, while in the other form, that of hijacking, the address may have been dormant for some time and its use is taken up by the hijacker. There are also cases where active addresses are being re-advertised incorrectly, either inadvertently, or as part of some form of malicious attack. All of these cases of unauthorized use of address resources fit within the broad term of a *bogon*. Sometimes a bogon is just a case of keystroke error by a network operator, and the consequent bogons are entirely inadvertent, and other times it may be a disagreement between an end user and a registration authority, and sometimes it may indeed be an instance of deliberate hijacking of an address.

A list of current possible bogons is published at the CIDR Report, among other sources. This report is updated on an hourly basis, and covers the use of IPv4 addresses and AS numbers where there does not appear to be any associated registration information.

The CIDR Report is an online resource relating to potential 'good housekeeping' actions within the inter-domain routing space. The reports are updated on an hourly basis. There are as number of reports, including aggregation and bogon reports.

Aggregation Report: http://www.cidr-report.org

Current Bogon Report: http://www.cidr-report.org/index.html#Bogons

Bogon Filters: http://www.cidr-report.org/bogons

What's the problem with Bogons?

Why should we worry about use of bogons in the Internet, and go the trouble of assembling lists of registered and unregistered number blocks? There are a number of reasons why admitting bogons into the Internet can be seen as a problem we should not take lightly.

The first reason is that in the address realm careful housekeeping is the first line of defense against attack. This is based on the observation that some attacks have used bogon addresses as their launch platform. This form of attack is not so common these days, as many, if not most, forms of attack now use enlisted 'zombies', where a number of systems are effectively taken over by the attacker and then used as part of a distributed denial of service attack. In this manner it appears to be relatively easy to enlist systems that have 'real' addresses to launch an attack, and there's no real need to inject bogons into the network as part of the attack. This leads to the question of why should we bother about bogons if the attack patterns have shifted away from bogons in any case? One of the counter arguments is that the attack patterns have shifted in recognition of the common awareness of bogons, and have shifted in response to the pattern of bogon filtering that exists in the Internet today. The concern is that if we pulled out all these bogon filters would we then see a return to the use of bogons in attacks?

The second reason is based around the observation that the process of obtaining address space involves disclosure of identity and cost, as well as meeting a number of policy criteria. If your objective is to use an address block for less than completely upright and honest reasons, and you really don't want to be readily identified, such as if you are thinking about sending a burst of unsolicited mail or considering other forms of malicious behavior, its tempting to steal an address block and launch a routing object into the Internet for a few days, and then move on. Again here the first line of defense against this form of abuse is being able to detect the activation of bogon address blocks and prevent their use within the Internet.

Of course there is also the bigger question of why we use registration processes at all, and the question of the role of the Regional Internet Registries (RIRs) in the first place. The answers to such questions lie in the guarantees of uniqueness and the uniform application of policy constraints that such processes support. If there was no registration process and no common policy, and if address allocation was a loosely coordinated (if at all) free for all, then its likely that we'd not have a useable network as a result. In its place we'd have chaos. What prevents a third party attempting to hijack your address block, or attempting to seize exclusive control over all the address space, or otherwise play fast and loose with address space is indeed this registration process and the associated policy framework. So most players abide by the policies and, nire generally, abide by the intent of these policies. Address uniqueness is the foundation upon which the network is constructed. Without uniqueness of addresses you have address chaos. And if you have chaos in addresses distribution function then the network itself falls. And for those who choose not to play by these common rules and accepted practices we have bogons and bogon filters.

What's a Bogon Again?

So we've noted that a bogon is an unauthorized use of an IP address. Fine words, but how do you enter this into your router as a filtering condition?

At this point the precise definition of a bogon becomes a matter of some interest, and in order to define a bogon its useful to define how an authorized use of an address is derived.

The first point of allocation is the Internet Assigned Numbers Authority (IANA). The registry of Addresses, and Autonomous System numbers can be found in a number of IANA registries.

> The Internet Assigned Numbers Authority operates a collection of protocol parameter registries, that record the distribution of protocol parameter values to third parties. In the area of number resource management the following registries record IANA's actions.
>
> IPv4 Address Registry:
> http://www.iana.org/assignments/ipv4-address-space
> IPv6 Address Registry:
> http://www.iana.org/assignments/ipv6-tla-assignments
> Autonomous System Number Registry:
> http://www.iana.org/assignments/as-numbers

These IANA registries describes which address blocks are available for use in the network, and which of these blocks have already been handed out for further distribution, and which have not.

So the first definition of a bogon includes all number blocks that are being held as 'reserved' by the IANA. These three lists (IPv4, IPv6, ASNs) can certainly be described in a compact format. The associated filters that can be generated from these lists can be used to filter inter-domain routing updates, to ensure that the local routing domain does not inadvertently learn a route to a bogon destination, or as packet filters, to ensure that the local domain does not admit a packet that is sourced from a bogon address.

Sounds good? So far so good. But the problem here is that when you look closely at the IANA registries there are a collection of errors and inconsistencies that, correctly, lead to questioning the validity and completeness of the data. In other words there appear to be errors and inconsistencies in the IANA data. Whoops!

> A list of the problem's I've encountered in attempting to use the IANA IPv4 address allocation data can be found at:
> http://www.cidr-report.org/bogons/iana-data.html

Uniqueness is a very fragile property of any space, and the foundation of any system of the distribution of unique tokens relies on consistent and accurate record keeping across the entire system. Its a problem, to say the least, when the root IANA registry is not completely accurate and completely unambiguous. Whoops indeed.

Assuming that these registry errors and anomalies are resolved in some fashion, then in IPv4 address space terms we've managed to identify the 31% of the IPv4 address space that is currently held in reserve by IANA. If we also include the 14% of the address space reserved by the IETF for various purposes, including the private use space, the bogon filter based on IANA data spans some 45% of the IPv4 space.

This kind of bogon filter, based on IANA data, has a role in deployed networks where it appears to server a role of catching inadvertent leakage of private use number space into the public domain. In the case of some AS number blocks it appears to also trap the use of historically allocated AS numbers where no registration information appears to exist. Unfortunately we really want to trap and prevent the first and not the second kind of use. The second type of use appears to be the outcome of an unresolved accident of record keeping rather than deliberate subversion of the integrity of the address space.

Bogon Beacons

Such an 'IANA bogon' filter needs to change every few months as IANA allocates new address blocks to the RIRs.

The problem that has been noticed is that these filters are often considered "set and forget" configuration objects. When the IANA allocates new address blocks to the RIRs there are a number of networks that do not update their bogon filters, and the unfortunate recipients of address space drawn from these blocks only get a partial view of the Internet.

Can we identify these filter laggards? It appears that a tool to test the propagation of routes through the network can be constructed by combining the concept of inter-domain routing update collectors and the deliberate advertising of prefixes from these address blocks. The folk at the RIPE NCC have done precisely that in their "bogon beacons". The intent is to provide some indication of where older packet and routing filters may exist that prevent such recently allocated address blocks from being seen by the entire Internet.

> De-Bogonising New Address Blocks: The following reference describes the approach used by the RIPE NCC to create a toolset that allows some form of detection of bogon filters that are triggering on recently allocated address blocks.
>
> http://www.ris.ripe.net/debogon/index.html

So at this point we have in our bogon inventory a definition of a bogon, a relatively simple filter set that can catch packets that attempt to use these bogon addresses, and an associated filter that can be used to catch routing protocol updates that attempt to install routes for these bogons. As well, we now have a way to test the coherency of recently-allocated address blocks, to detect instances where bogon filters have not been updated to reflect the most recent IANA allocations.

Good enough? Unfortunately not! Not because the technique is bad - far from it. The bogon beacons are a reasonable approach to the problem of identifying the location of poorly maintained bogon filters. The problem lies more in the working definition of a bogon, which is still incomplete.

More Bogons

If you are attempting to hijack some address space and you'd like to get it routed by everyone else, then using address space reserved by IANA is probably too obvious. Perhaps it would be more devious to use a block of address space that has already been allocated by the IANA, so that it would slip under our original bogon filter. But if we don't want to alert the any user of 'real' address space, then maybe what we are looking for is address space that has been allocated by the IANA into the RIR system, but the RIRs have not yet allocated the address space to the customer. This is a relatively large block of address space, spanning some 9% of the total IPv4 address space at the time of writing this article.

The data that we can use to base this determination of as-yet- unallocated address space comes from the RIR's own allocation records. Generally, this data appears to be consistent and trustworthy. Generally. Unfortunately there are some inconsistencies in the details, and again some level of judgment is called for to fully resolve this data.

> The Regional Internet Registries publish a record of all allocated number resources. These files are updated daily to reflect the current allocation status. There is also a set of older files on a daily basis
>
> APNIC: ftp://ftp.apnic.net/pub/stats/apnic/delegated-apnic-latest/
>
> ARIN: ftp://ftp.arin.net/pub/stats/arin/delegated-arin-latest/
>
> RIPE: ftp://ftp.ripe.net/pub/stats/ripencc/delegated-ripencc-latest/
>
> LACNIC: ftp://ftp.ripe.net/pub/stats/lacnic/delegated-lacnic-latest/

At this point the bogon set encompasses the equivalent of 137 /8 address blocks, but its heavily fragmented into small 'holes' that complement allocations. Almost one half of this address space reflects address space that was originally allocated in the 1980's and 1990's, and the space reflects piecemeal returns back to the registries. If we wanted to make a comprehensive filter list of bogons, encompassing the address space that the RIRs currently hold and have not yet allocated for use, then the most compact form of the filter list is now some 7,200 entries (http://www.cidr-report.org/bogons/freespace-prefix.txt). And, the list changes every day, so you will probably need to use a script that looks for new allocations from the RIRs and updates all the bogon filters on all routers on a continual basis.

This type of bogon list is getting too big to be installed easily on routers, and, as it changes in detail each day as each RIR performs further allocations, there is a requirement to set up elaborate scripts that operate a frequent intervals. The space is also heavily fragmented, so the concept of bogon beacons is no longer all that useful. Its also the case that there are some poor or non-existent records of allocated resources, and some allocations appear to be the subject of dispute between the RIR and the entity currently using the resource. Other sources of data, such as the 'whois' databases operated by the RIRs can tell a different story about the status of address blocks. So as we attempt to get a more accurate view of bogons, we also expose some additional uncertainties in the data at this level of detail (http://www.cidr-report.org/bogons/rir-data.html).

Assuming we could resolve all these inconsistencies and obtain an up- to-date and accurate real time feed of all currently allocated address blocks, which by any reckoning would require a considerable additional effort over the current situation of IANA-based filters, there is still a question of completeness.

With all this effort would we have a useful and complete list of all possible bogons? Well if we want to encompass all possible forms of address hijacking, then once more the answer is, unfortunately, not.

There is a further 16% of the address space, or the equivalent of 42 /8 address blocks that have been passed from IANA to the RIRs as part of normal allocations, and the RIRs have allocated the addresses to end customers, but there is no trace of these addresses in the routing table. This is a lot of address space, and while up to date allocation information is held for much of this address space, unfortunately the allocation information is incomplete or missing for some of these address blocks. Its these latter blocks that are perhaps the most tempting target for address hijacker.

And unfortunately its impossible to create a filter set that can mask out these addresses, and the criteria for selecting these addresses as potential hijacking targets is highly subjective.

So filtering appears to be easiest when its at its most ineffective. The IANA filters are easy to define, but are woefully incomplete. Despite their shortcomings as a meaningful bogon filter they are used more frequently than is probably merited. More detailed filters are more challenging to maintain, but even those filters are incomplete and ineffectual at isolating determined instances of hijacking of address space. A complete set of potential bogons can be assembled, but its extremely challenging to script up a means of updating these filters on very frequent basis.

Maybe filters are just another form of response to network abuse that makes you feel a bit better that you are doing 'something', but in reality the outcomes are largely symbolic rather than functionally effective.

Alternatives to Filtering?

Perhaps filtering is not the answer. A number of recent efforts have tried to centralize the effort associated with generation of a bogon filter, and distribute the equivalent of a filter though BGP itself.

The idea is that a filter set gets assembled at one point as a collection of route objects with next hop addresses that point to a null route. Edge routers can use an eBGP multi-hop peering session to obtain a set of routes for these filtered addresses, and use local prefs to ensure that the null route is preferred.

It sounds tempting as a solution, and certainly the one compelling aspect of this solution is that there are not hundreds of different folk each attempting to place an individual interpretation on the various inconsistencies in the data. However its only a partial solution, in that it only traps certain types of bogon routing and certain types of packet forwarding. It does not trap the situation of the use of more specific routing objects, nor does it trap the use of bogon addresses as source addresses in packets, nor does it trap the use of bogon ASNs in the AS path. But on the theory that one partial solution is as good as any other, this particular partial solution looks like a whole lot less work.

What's the real problem?

I suppose the problem with filters is that it appears to be yet another case of the general observation that when all you have is a hammer, everything looks like a nail. The problem relating to bogons in the Internet can be stated more generally as a lack of 'good' information about what are 'valid' or 'authorized' addresses that we should see in routing exchanges and that we should see as source or destination fields in packets.

The approach we've seen to address this is to install checkpoints along various network paths (filters) where routing objects and packets are subject to scrutiny.

Of course there are other ways to achieve this overall objective, and one way is to make the entire test of authenticity more overt. This would entail a registration entity being able to attest that a particular address block is managed by a named entity. Through use of digital certificates, the entity could provide information to its neighbor network domains that it was the valid originator of an address route object. These neighbors could then counter-sign this certificate as the route object was further promulgated across the network. When a route object arrives at a remote domain the chain of trust associated with the AS path of the received route can be verified, as can the authenticity of the original injection of the route. Unauthorized use of an address block, or a bogon, would be identified by not having a reliable attestation of validity that is signed by the relevant RIR.

So the real problem not the problem with twiddling with filters, but its the lack of a reliable foundation of testable authenticity within the inter-domain routing environment. If valid routes had associated digital certificates that ultimately referenced the original assignment transaction, we would be more confident in asserting that the routing system contains only those objects where the owner of the private key of the assignment entity has approved the original injection of the route. We might not stop the entire use of bogons in the Internet, but within this system the attackers would need to undertake a complete identity theft in order to be able to inject the bogon routing object.

And what should we be doing?

This paints a picture that indicated that various forms of filtering as a means of bogon prevention are really not all that effective, in that they are incomplete in content, incomplete in deployment and rely on a level of guesswork about what is authentic in terms of use of addresses.

The alternative appears to be that we make authenticity of a route object a testable assertion and for this we need to place the registration entity in the role of being a trust point, with the original injection of a route and the AS path of the route object advertisement being actions that others may not accept until they have independently verified these actions through checking the validity of the associated digital certificates.

Securing inter-domain routing appears to be the real agenda here, and some form of secure BGP appears to be much in need.

There are two approaches to secure BGP currently published. A long-standing effort is that of "secure BGP", developed largely in the Internetwork Research area of BBN Technologies.

http://www.net-tech.bbn.com/sbgp/sbgp-index.html

Another approach, developed within Cisco, is "Secure Origin BGP". This is documented in an Internet draft:

draft-ng-sobgp-bgp-extensions

Both approaches have been described in the Internet Protocol Journal, Volume 6, Number 3, Sept. 2003. Its worth the read, both to gain some understanding of the basic concepts behind these proposals, and also to compare the two approaches in terms of trust models and functionality.

http://www.cisco.com/warp/public/759/ipj_6-3/ipj_6-3.html

I would've thought that it should be on the critical action path for ISPs to be looking at the topic of securing inter-domain routing in all its detail, and then assisting their local RIR to undertake the role of trust point for the allocations they manage. I would've thought that the IETF would be looking at technology proposals in this area and undertaking a serious effort to understand the relative strengths and weaknesses of the proposals and then working diligently towards a single mechanism that could secure the inter-domain routing space. I would've thought that customers should be sending clear signals that they want and need secure BGP. I also would've thought vendors would want to provide protocol implementations that conform to well considered standards and are interoperable and robust. I would've through that the RIRs themselves would be trolling through the various databases to get a clear and coherent history of address allocations, so that if they are asked to provide an attestation that a certain entity has control over an address block it can counter-sign this attestation using its own data sources.

In this case the current scorecard is not much different from something slightly less than one out of three. There is some visible work within the RIRs to improve the accuracy and consistency of their data, and work related to lay the foundations of the issuing of digital certifications, which is really encouraging, as its an essential piece of groundwork for this activity. But that's not enough by itself. If we really want to go bogon hunting we should be doing the rest of the job as well.

If we care about the usefulness of the Internet then we should care deeply about the integrity of its address space and its use in the Internet. And if this is something we care about then we should be supporting the integration of explicit authentication of route objects in the routing protocol itself. We can eliminate bogons, but not though anything less

than an concerted effort to properly secure inter-domain routing. Bogon filters are simply a distraction from the critical agenda here.

Geoff Huston

Disclaimer

The above views do not represent the views of the Internet Society, nor do they represent the views of the author's employer, the Telstra Corporation. They were possibly the opinions of the author at the time of writing this article, but things always change, including the author's opinions!

About the Author

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He has been closely involved with the development of the Internet for the past decade, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. Huston is currently the Chief Scientist in the Internet area for Telstra. He is also the Executive Director of the Internet Architecture Board, and is a member of the APNIC Executive Committee. He was an inaugural Trustee of the Internet Society, and served as Secretary of the Board of Trustees from 1993 until 2001, with a term of service as chair of the Board of Trustees in 1999 and 2000. He is author of a number of Internet-related books.