**The ISP Column**
*A column on various things Internet*

July 2018
Geoff Huston

# An Update on Securing BGP from IETF 102

One way or another we've been working on various aspects of securing the Internet's inter-domain routing system for many years. I recall presentations dating back to the late '90's that point vaguely to using some form of digital signature on BGP updates that would allow a BGP speaker to assure themselves as to the veracity of a route advertisement. The concept is by no means a new one, and even the approach of digital signatures has been part of the conversation since its inception, as shown in an industry presentation from 1999.

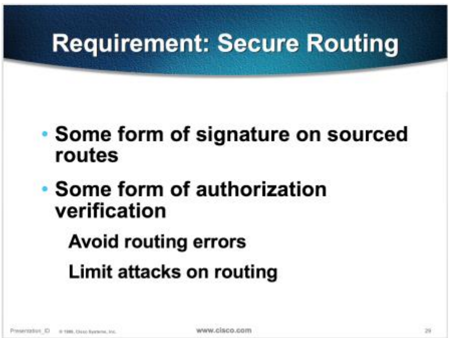

**Requirement: Secure Routing**

- **Some form of signature on sourced routes**
- **Some form of authorization verification**
  - **Avoid routing errors**
  - **Limit attacks on routing**

www.cisco.com

*Excerpt from a Fred Baker and Dave Meyer presentation on "Global Routing Issues", Cisco, 1999*

Some twenty years later we are still working on securing the routing system.

In this article I'd like to look through some items that have come up during the July 2018 meeting of the Internet Engineering Task Force (IETF) and try to place these items into some bigger context of routing security.

## A Potted History of Securing BGP

However, before we start that, it might help to give at least a rapid-fire summary of what has happened along this path to a secure routing system for the Internet.

A very early step along this path was the *whois* registry that was operated by the network registry of that time, InterNIC. You may not understand why something odd was happening in the routing system, but if the network was still working you could use a really simple query protocol and get the identity of the operator of a domain name, an IP address or Autonomous System Number, and maybe even a telephone number! It may not stop various forms of fat finger slip-ups or various deliberate efforts to subvert the routing system but at least you could call someone who was listed in the whois database as being associated with the IP address.

The next major development was that of the use of Route Registries and the associated Routing Policy Specification Language (RPSL). This was an outcome of the routing arbiter projects of the early 90's. This allows a query agent to look up an IP address prefix and retrieve the routing policy intent of the address' owner. RPSL is a formal language that is intended to be machine parsed, and one objective was that network operators could use this route registry information to construct filters that admitted only what the address-originating network operator had said they intended and deny all else. This model of allows a network operator to define their routing practices in advance and then allow others to link to these registries and automatically filter what is actually advertised to ensure that it remains consistent with their stated intentions. Route Registries are still around today, and have been very useful in many ways, but they are by no means a panacea. These days there are many, probably too many, routing registries and the sum of all information contained in them can be mutually inconsistent, which leads to considerable confusion. There is no clear authority model that would allow a registry client to determine what data can be trusted, and little in the way of active curating of the routing data to ensure that it is current, consistent and useful. It appears that routing registries look like a very sensible approach in theory, and they have worked effectively within certain communities. However these are isolated successes, and it appears that our more general experience has been relatively uninspiring so far.

There have been a number of efforts concentrated on securing the BGP routing protocol. BGP used unencrypted TCP sessions and was vulnerable to man-in-the-middle attacks. Securing these sessions prevents efforts to tamper with the BGP information in flight. Of course, part of the issue here is that within the 62,000 networks that collectively comprise the Internet, it would be unreasonable to claim that all of these networks are eminently trustworthy. If a hostile actor can control a BGP speaker then it is possible to inject falsified information into the routing system despite channel protection. So, it appears that while protecting the channel is a good move, it's just not enough, and we need to look at measures that can protect the routing content as well.

One of the ways we can protect the integrity of the routing content is to use cryptographically signed attestations in routing. Protecting BGP updates with digital signatures allows a BGP speaker to make a judgement as to whether a received advertisement represents the original intent of the owner of the advertised address block and that there has been no untoward tampering of the routing information while it has propagated over the inter-AS routing space. The underlying approach binds cryptographic keys to address blocks in a robust manner (using X.509 public key certificates in a simple hierarchy) and use these keys to sign a digital contract to permit the advertisement of the prefix. A Resource Public Key Infrastructure (RPKI) allows these digital signatures to be validated, and the RPKI itself binds a holder of a key to the role of controller of an allocated IP address. Origination information can be digitally signed by the address holder as a proof of authority, and this can be coupled with a means of signing over the AS-Path attribute of a BGP Update, so that a receiver of the update can derive some confidence that this part of the update has not been tampered with in any untoward manner.

There have been a number of proposals that use systems other than X.509 certificates to bind the holder of a public key to the role of the controller of a block of IP addresses. Most recently we've seen some thoughts about applying Blockchain technology to this space. A blockchain provides a means of associating a cryptographic key with an IP address block and wrapping it with a digital contract. The digital contract can be used to express a routing permission, and, presumably, could also be used in some manner to carry information relating to an inter-AS routing relationship.

Which just about brings us to the present with securing BGP.

## What did we hear at IETF 102?

### RPKI Validation Reconsidered

The first item is a rather obscure discussion about the way we could implement a change to the validation procedure for these RPKI certificates. (https://datatracker.ietf.org/meeting/102/materials/slides-102-sidrops-deploying-validation-reconsidered-00)

Some years ago, it was proposed to make a change in the procedure used to check if a RPKI certificate was indeed valid. The change was intended to remove some of the brittleness of coordination of certificates between various parties and allow certificates to stray slightly from a strict encompassing relationship.

I just can't bring myself to spend too long in the details here, as they are less of an issue than the consequences. The way the change was adopted in the IETF standards was a change in the certificate identifying code (the OID) such that these certificates that specify that they may be validated with the new validation algorithm are not backward compatible with the older certificates. With me so far?

If we can't mix and match these OIDs in a validation path, then we have a problem. Either every certificate gets published twice, once with the old OID and once with the new OID, which seems to be just crazy, or we have a ‘flag day’ and on that day the old OIDs are withdrawn and the certificates with the new OID are published.

These days the concept of a flag day is about as popular as the plague. The Internet is just too big for a flag day for most things (think IPv4 to IPv6 transition) and while some are of the view that the RPKI user community is still small enough to contemplate a coordinated flag day, others were arguing that even this moment has passed, and a flag day will cause an unacceptable level of disruption. This seems a lot like getting stuck between a rock and a hard place.

Doing nothing has its own drawbacks, as there are concerns that the existing framework is overly brittle. This is cited as a reason why some network operators are just not interested in deploying the solution. But if we cannot propagate changes to the RPKI that appear to improve its operational robustness, then the chances of widespread, indeed universal, deployment look slim. And why is this important? Because like many security systems, this system can only sign what is ‘good’. There is no 'this is bad” signature. In the absence of an explicit evil bit we rely on a more basic observation: if all the ‘good’ objects are validly signed then all else is therefore bad.

Personally, I believe that it’s still very early days for the RPKI and arguing that we cannot make these changes to the system in these early days because of the existing deployed base is too large to change only strengthens the case that the RPKI’s ultimate fate is to be one more tombstone along a long and winding BGP Security path. Crypto systems require operational care and precision and are harshly intolerant of any level of aberration. They require accurate time, proper key management, due attention to robust availability of published material and a readiness to react quickly and intelligently to the slings and arrows of the inevitable operational mishaps. Many operators see this entire exercise as one more area of operational fragility and are justifiably cautious of such systems. There is value in attempting to increase the robustness of these systems if we want the RPKI to have some chance of success.

## ROA Maxlen Vulnerabilities

The entire package of BGP security had four components: a PKI of binding public keys to IP address blocks, a structure of the publication of these public key certificates in a distributed repository and the associated collection of this distributed data into a local complete corpus, signed attestations that authorise a network to advertise reachability for an IP address block and finally a method that allows routers to sign across the AS Path attribute of BGP update messages to protect the AS Path from tampering.

This last component, namely the BGPSEC protocol, is the one that looks the most unlikely to be universally deployed at this point in time, yet the inability of the protocol to ‘bridge’ between islands of deployment does infer that it is only really useful if it is universally deployed! If the realistic prospect for this BGP security ‘package’ is partial deployment, then what can be salvaged from this that still has benefit even when only deployed on a piecemeal basis?

One presentation in the SIDR Operations Working Group looked at the use of ROAs and proposed a simple measure that would reduce vulnerability to some particular form of address hijack. (https://datatracker.ietf.org/meeting/102/materials/slides-102-sidrops-draft-ieft-sidrops-rpkimaxlen-00) In the early days of the development of the ROA, the ROA simply listed a collection of IP address prefixes and a collection of ASs that were authorised to originate routes for these addresses. In an envisaged world of universal deployment any advertisement not described by a ROA was immediately suspect. However, the ROA semantics were altered to specify a single address prefix and add a ‘maxLength’ attribute. The semantics of this maxLength field were somewhat widespread, and best explained by example.

> If a ROA contained 10.1.0.0/16, maxLength=24, AS 3 then AS3 was authorised to advertise 10.1.0.0/16 as an originated route. But as well it could advertise any more specific address prefix up to a /24, so 10.1.1.0/24 was also valid, as was 10.1.2.0/20. But a ROA also has an implicit claim of what is

invalid. Any advertisement of length more than maxLength was invalid, and any other AS originating these routes was invalid, unless of course another valid ROA described these advertisements.

All this looks quite reasonable until you think about a world there are only ROAs and there is no BGPSEC AS Path protection. How do you hijack a prefix? The easiest way is to advertise a more specific route.

Let's now take the case where an address holder has minted a this 10.1.0.0/16 maxLength=24, AS 3 ROA but only advertises the /16. A BGP hijacker could advertise both 10.1.0.0/17, origin AS 3 and 10.1.128.0/17 with a faked origin AS 3 and the ROA would make these fake route advertisements look legitimate. The two covering more specifics span the aggregate, so all addresses are hijacked with this approach.

Obviously AS Path protection could prevent the faked insertion of the origin AS, but without AS Path protection all that's left is to warn ROA publishers not to be overly permissive with the maxLength field of ROAs. This will not stop a potential hijack of an address, but the attacker is then unable to claim the entire address by using more specific prefixes. Instead, the best an attacker can do is to perform a partial hijack based on route propagation and AS path length preferences when the genuine and the hijack advertisements specify the same address prefix. Perhaps the ROA maxLength was in retrospect a very poor idea.

### AS Adjacency Attestations

It seems uncomfortable to just give up on AS Path protection completely, and we are left wondering if we can salvage something from this.

Back around 2000 there was a second proposal to secure BGP, namely secure origin BGP (soBGP). It was devised before the RPKI matured, so many aspects of its derivation of trust look somewhat weak in retrospect, but one aspect in particular still looks attractive today, namely the AS adjacency attestation.

What if each network operator generated a set of AS Adjacency attentions that enumerated all the AS's that their network had a routing relationship? If all AS's did this, then a BGP speaker could test the AS Path of a received update against the collection of pairwise AS adjacencies and reject the update on the basis of a faked AS Path if there was an AS pair not described by an AS adjacency attestation. This implies that a potential BGP hijacker could only create a faked AS path using existing attested AS adjacencies. In other words, the faked AS path would be a plausible AS path in the first place, and the freedom of a hijacker to invent synthetic shorter AS paths that attract traffic through some man-in-the-middle is severely curtailed.

This approach does not require BGP speakers to sign BGP update messages or validate them at the other end. Like ROA processing, the crypto component can be offloaded, and a set of AS path constraints passed to BGP speakers as either a filter or a BGP black hole feed.

However, universal deployment of such a mechanism is unlikely, so what is the benefit of partial deployment of AS adjacency attestations? If the condition is placed on AS holders that if they mint an AS adjacency attestation about one AS adjacency then they mint an attestation for all adjacent AS's, then this can be useful even in a partial adoption scenario. If a route hijack uses a forged AS Path then if the hijacker includes an attestation publishing AS in its forged path, even as an originating AS, then it must also include one of the adjacent AS's as listed in the adjacency attention. The higher the level of adoption of adjacency attentions the more challenging it is for an attacker to forge a plausible AS path.

This AS adjacency model was revived in an internet draft back in 2010 (https://tools.ietf.org/html/draft-huston-sidr-aao-profile-03) but the work was abandoned at the time due to lack of interest. The SIDR OPS session at IETF 102 saw a further reprise of this work (https://datatracker.ietf.org/meeting/102/materials/slides-102-sidrops-as-path-verifcation-using-aspa-00).

This latest approach to AS pair adjacency attestation refines the earlier approach by defining only a single AS adjacency pair and removes the constraint of explicit enumeration of all an AS's adjacencies. The draft adds a policy component by defining the adjacencies described in this manner in a Customer / Provider relationship.

It remains an open issue how much publication of routing policy information in BGP is sufficient and how much is too much, to the extent that it deters AS's from publishing any information at all. Whether or not policy is included, if the constraint of complete enumeration is retained then there is some value for an AS in publishing these adjacency attentions even in a partial deployment scenario. For those prefixes the AS originates a potential hijacker has to not only use the same origin AS to satisfy the ROA constraints, it also has to include the first hop AS in the synthetic AS Path. If this first hop AS also publishes its attestations the attacker also has to include the second hop AS, and so on. An attacker cannot lie about the adjacencies of an AS if the AS has published these adjacencies.

## soBGP

It has been an interesting exercise to compare the current state of SIDR with soBGP from about a decade ago. While SIDR OPS is ostensibly an operations working group it has been completely unable to resist the temptation to delve back into the architecture and protocols of secure routing systems and in so doing the the work appears to be reproducing quite accurately all of the salient elements of soBGP. Parenthetically, it is somewhat disappointing to see this reuse of the earlier work without any due credit, as this seems entirely foreign to the IETF's usual practices of careful attribution.

But as we turn to soBGP as the way to salvage some viable means of securing the operation of BGP, I suppose the real question is why did we abandon soBGP before? If we are rediscovering its positive attributes can we also try and remember its negatives?

As I recall, the argument at the time was one of the distinction between AS Path *plausibility* and AS path *provability*. The soBGP protocol does not attempt to validate that the update actually traversed the path that is described by the AS Path attribute. It may be a synthetic path and soBGP cannot detect that. But what soBGP can reveal is whether the path, or elements of the path in partial deployment scenarios, is a *plausible* path. BGPSEC imposes a far stricter condition, namely that the update traversed the AS path exactly as described in the AS Path attribute of the update.

But this stricter provability condition is only available when every BGP speaker operates BGPSEC, when every eBGP router is provisioned with keys and every eBGP speaker signs announced routes and validates received updates. Perhaps we were seduced by the prospect of a highly automated secure system and it was only later that it became obvious that BGPSEC has too many deployment impediments and universal deployment (a pre-requsite for BGPSEC) is simply unachievable. A revival of soBGP can make use of the RPKI and ROA infrastructure, and can provide a workable outcome by simply adding the component of AS adjacencies to the mix. Critically, partial deployment of AS adjacency attestations provides additional levels of assurance to the networks that publish these instruments, in that attempts to hijack prefixes originated by these AS's or use them in a forged AS Path attribute require the inclusion of adjacent AS's, further lengthening the AS Path length of the forged path and reducing the effectiveness of the attempted hijack.

## Blockchain

These days Blockchain has certainly achieved mania status and it seems that any researcher seeking funding is now obliged to add the magic word *Blockchain* into the proposal, irrespective of the subject matter of the research itself!

It comes as absolutely no surprise to see what we might call "BGPcoin" appearing in this space as a potential response as to how to secure BGP. The essential characteristic of this approach is a distributed ledger that operates without a governing authority to manage the ledger. It is based on the formation of a consensus that is recorded in a public tamper-proof transaction log. Demonstration of control or ownership of a resource is based on the association of a description of the resource and the public key of the owner, placed into a block that is inside a Merkle chain.

In the DINRG a research group presented their take on this approach (https://datatracker.ietf.org/meeting/102/materials/slides-102-dinrg-decentralized-internet-resource-trust-infrastructure-bingyang-liu-00). One element they are using is a so-called "smart contract" which represents a unilateral assertion that a public key holder owns an IP address block for the period of the contract. The contractual process engine draws an address prefix from the available address pool and enters it into the ledger. The entry is incorporated in a block of transactions, and the Merkle chain binds the contents of the block against the previous block in an immutable manner. The other element is the mapping of the BGPSEC ROA object, where a transaction is the association of an IP address prefix with an AS number.

This particular blockchain approach replaces the address registry operators, the secure credential system of the RPKI and the attestations of the ROAs in the blockchain. Before we all head off in a wild mania of excitement that bitcoin technology solves everything there are some downsides here that are unclear. If a party loses access to its private key or does not renew the address ownership contract it has no further access to this address and if the address is being used in a network then the network has to be renumbered to a new address. It is unclear to what extent the conventional registry contents of the description of the identity of the holding entity is incorporated into the ledger entry. Validation of individual ledger entries requires access to the entire blockchain, and bloating of this data set is a risk. The process of consensus is also fragile, and in other research work presentations at this ,meeting we heard of attack vectors that cause consensus to split and the blockchain to bifurcate. Without a central authority to rectify this situation it is unclear how the ledger is returned to a single collection. There is also the risk of takeover, in that if an attacker can control more than one half of the entities that actively form consensus of the block chain then it appears to be possible to subvert the ledger.

As the internet continues to grow, and as it assumes a central role in a global society driven by data and digital content, we naturally treat the network in a more conservative manner. Experimentation with some of these critical elements of the Internet, including the role of name and address registries, and experiments with novel approaches to secure these elements no longer excite a reaction of "let just try it out and see how it works". These days we are far more critical of potential changes and we are far more conservative in our assessment of risk when we consider such experiments. With that in mind, these decentralised distributed systems have to demonstrate a greater level of robustness and acceptance prior to any form of adoption in the Internet.

## The State of Securing BGP

My intuition is that we are still quite comfortable with the existing human-operated central name and address registries, and it would take some crisis of confidence to adopt a completely different model. Given that consideration, it is likely that the RPKI will also persist for some time, and it looks likely that it may pick up greater levels of adoption across the Internet. However, prognostications about universal adoption remains far more challenging to phrase.

It is also likely that ROAs are sufficiently useful to remain a part of the picture, but the prospects for BGPSEC are looking pretty poor. Of course, without some form of protection of the AS Path route hijacks are relatively easy to mount by simply including the ROA-specific origin in the attack. And with liberal use of maxLength in the ROA more specific attacks will be effective.

We are back to a question that was considered at length and without any conclusion by the earlier RPSEC Working Group in the IETF some years ago. How can we protect the AS Path in BGP? Having tried and evidently failed with a relatively comprehensive approach in BGPSEC, it appears that we are back to the more pragmatic approach of soBGP and AS Path plausibility checks. The current proposal for AS adjacency that attempt to unite elements of routing policy with inter-AS topology may well be attempting to do too much. One area of consideration at this point in time is whether it would be better to maintain a clear distinction between the role of securing the topology part of BGP that maintains a view of the connected Internet at an AS level, and the role of securing the policy framework of routing that imposes a selective filter on routing information to determine how data paths are maintained within this topology.

There is one thing we do know about the effort intended to secure the Internet's distributed routing system. It's taken some thirty years of study and experimentation without arriving at a clear and obvious solution because it is a very challenging problem. It would be nice to think that we are coalescing to a practical set of measures that create a pragmatic response to route hijacking. But if you think I'm being unduly optimistic in such a call, then I can understand that many commentators hold a more sceptical view of this effort. Irrespective of whether the work is coalescing or not, it's obvious that interest in this tropic is not waning, and we'll see a lot more discussion of this topic at IETF 103.

## Disclaimer

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

## 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*