

Securing Inter-Domain Routing

While a network administrator can do a relatively robust job in protecting their local routing infrastructure by taking reasonable measures in terms of protecting routers and protecting the integrity of the operation of routing protocols within their local network, the task of protecting the integrity and accuracy of the information being carried by the routing protocols is a far more challenging task. The current level of activity, while well-intentioned, is not proving to be resilient against even such benign forms of routing information corruption as operation misconfiguration, let alone being resilient against determined and creative attack vectors.

It's an area where the rewards of mounting a successful attack in the routing system can be very high. It is possible to undertake denial of service, third party traffic inspection and service cloning, and to do so in a manner that can be challenging to detect through deliberate corruption of the information carried in the routing system. For example, by injecting a false route describing a path to an anycast instance of a root server, an attacker can create a network sub-domain where DNS resolution queries are passed to a fake server, who, in turn and manipulate the responses associated with the intended victim, while answering all other queries accurately. Or, more directly, an attacker can create a fake version of an online commerce service, subvert a sub-domain of the internet with false routing information and thereby harvest users' access credentials. Obviously the rewards of attack can be high, whether it's a targeted attack against a single service or a coordinated effort to disrupt the operation of large sections of the network.

The potential outcomes of attack on the routing system can include:

- **Blackholing** – where false routing advertisements redirect traffic away from its intended destination and instead are directed to a sink point. The outcome of this attack is an effective denial of service attack, where the target service is taken offline. A side-effect may be a rearrangement of traffic flows that could overload some network links.
- **Impersonation** – where false routing advertisements redirect traffic away from the intended destination and instead direct traffic to a site that masquerades as the destination service. Commonly this form of masquerading is used to gather otherwise confidential information from users of the original service.
- **Inspection and Alteration** – where targeting false routing advertisements cause traffic to an intended destination to be forwarded towards a compromised network segment, where the traffic may be inspected, or even altered before being passed onward to the actual destination.
- **Denial of service and Network Destabilization** – large scale generation of updates and withdrawals of routes cause severe capacity pressures on routers, to the point where routers may fail, which, in turn generates further instability in the routing system.

The attack mechanisms fall in three distinct categories:

- **Injection** – where false updates or withdrawals are generated and passed into the routing system.
- **Destabilization** – where routing updates and withdrawals are generated at a high frequency intended to trigger some form of route dampening response.
- **Overwhelming** – where a large number of updates are generated with the intent of exhausting the memory capacity of neighbouring routers and thereby causing routing failure.

In this article we'll look primarily at the issue of injection, and means to prevent third party passing off of injecting false information into the inter-domain routing space. It is admitted at the outset that this does not directly address the issues related to destabilization and overwhelming attacks, but the implication is that such attacks would need to manipulate verifiable routing advertisements, rather than using artificially generated data.

Routing Security

In the previous article we looked at some “good housekeeping” practices that are intended to prevent the local routing infrastructure from being subverted. After all, routers represent one of the more vulnerable points of weakness in the entire framework of routing security, and if an attacker can gain control of a router within a network, or gain control of its configuration, then using this platform to inject false information into the routing system is a logical next step. After all the collection of routers represent a trust domain, and it is normally the case that the collection of interior routers within a network are configured in a framework of mutual trust. So the first element of the security framework is that of protection of the routing elements themselves.

The previous article also looked at various ways to protect the integrity of operation of the routing protocols. The vulnerability here is that if a third party can inject packets into the routing protocol exchange, then, at the very least the attacker can disrupt the operation of the routing protocol and cause various forms of disruption and denial of service. There is also the potential to hijack a routing peering session and inject false information into the routing system. So the second element of the security framework is that of protection of the routing protocols.

The third element is to protect the integrity of the information being passed through the routing system. Here the current story is much weaker, in that there are few tools available today that can protect the integrity of the information being passed through the routing system. The basic questions that needs to be answered in a secure routing framework include establishing the bona fides of the party that originally injected the information into the routing system, as well as the bona fides of the routing prefix itself. This identity information is not of any intrinsic value in itself, but a means to answer the more fundamental question of whether this party has the necessary credentials or permission to inject this information. Have they been authorized to originate a reachability advertisement for this particular address prefix? The associated question is: Is the address prefix valid? In other words, has this prefix been duly allocated for use through the established address distribution procedures, and is it valid to use this particular address range in the context of the public Internet? And of course routing is used to support forwarding, so the associated question to be asked is: Is the forwarding path that is claimed to reach this destination one that reflects an authorized sequence of transit operations? This question is one that attempts to not only establish whether the forward path is valid in terms of connectivity, but also to establish at each point in the path sequence whether the transit actions conform to local policy and authorization.

What we have today is a routing framework that is incapable of answering these questions with any degree of authenticity.

Obviously this is less than comforting, in that the Internet’s routing system, one of the critical components of the overall framework of the integrity of operation of the Internet, appears to be loosely bound together by somewhat imperfect trust models rather than by any form of strong authentication. Basic information about the validity of address prefixes and autonomous system numbers is missing. Configuration errors are commonplace, and the consequences of such configuration errors range from mildly annoying to highly disruptive. The routing systems has already been the subject of attack, and it is assumed that the attacks will escalate in severity and frequency.

Route Filters

How can an ISP protect itself against such efforts to introduce false information into the routing system?

Some ISPs deploy route filters to attempt to set some limits on the routing information that will be accepted from a routing peer. These filters may simply encompass the private-use address space (the so-called “1918 filters”, named after RFC 1918, where the private use address prefixes are defined). A slightly larger set of filters encompasses the unallocated pool of addresses, and includes those address blocks that constitute the unallocated address pool that is administered by the Internet Assigned Numbers Authority (IANA). Such filters are often referred to as “bogon filters”. In theory it is possible to go one step further and create a fine-grained set of filters based on the RIR’s daily report files of assignment actions, although this is not a common form of filtering as used by ISPs. All three of these forms of filters can be used in the context of a routing filter, a packet filter, or both. However, they are coarse filters at the best of times, and can only act as a deterrent in deflecting traffic that purports to originate from reserved or unallocated addresses. Such filters do little to protect the local network from most forms of deliberate or inadvertent corruption of routing information.

The routing system today contains various forms of false information, ranging from advertised addresses that have no visible record of allocation, where filters have some role, through to hijacking of allocated address space and passing off at the service level. While many cases of such false information have origins in terms of historical record keeping or misconfiguration, a small, but nevertheless disturbing amount of this false information in the routing system is the outcome of deliberate attack.

So can filters assist in preventing the hijacking of allocated address space? Normally this is not the case, in that filters can identify obviously erroneous information, but they are not so effective in identifying those cases where the information refers to validly routeable address space, but where the routing information is not valid. In this case we probably need to ask a harder set of questions, of the form: Is the address space valid? Has the holder of the address space authorized the address space to be advertised by the holder of the originating AS? Does the routing path for this address prefix accurately reflect the likely forwarding path for traffic destined to this address prefix? If we want to answer these questions then we need to look beyond filters and examine BGP as a protocol, and then look at means to secure BGP routing.

BGP – The Protocol

BGP operates as a reliable two party message passing protocol. BGP uses TCP as its transport protocol, eliminating the need to implement message fragmentation, retransmission, acknowledgment, flow control and sequencing within the protocol itself. The BGP speakers exchange Update Messages. These messages have three parts: a BGP Header, a (possibly null) list of withdrawals, corresponding to those prefixes where the peer BGP neighbour no longer has a valid path, and a set of updates where the peer BGP has installed a more preferred route, or has learned the route of a new prefix. The update section is split into two parts: the common AS path attributes, and the set of address prefixes. This is indicated in Figure 1.

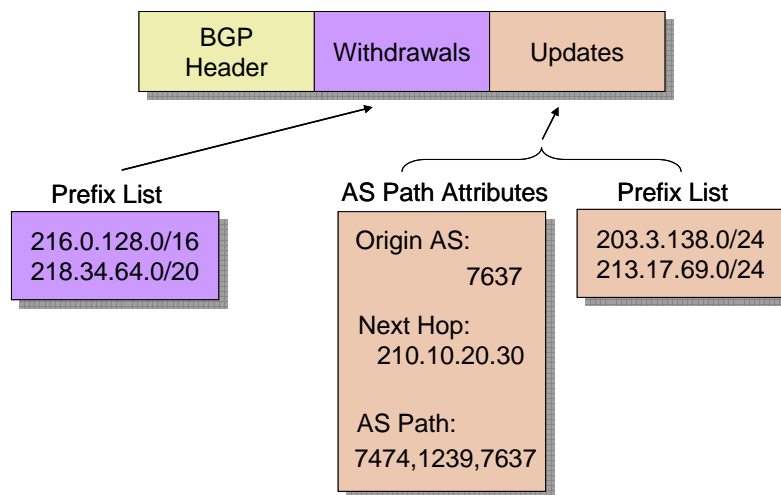


Figure 1 BGP Update Message Format

In processing this update message, the local BGP speaker updates a local adjacency state that is being maintained with the BGP peer, removing withdrawal routes, and adding or modifying the updated routes. The changes to this adjacency state (or “Adj-RIBs-In” to draw from the terminology used in conceptual BGP model as described in draft-ietf-idr-bgp4) are then processed by the BGP routing algorithm, which uses the BGP route preference rules, coupled with local policy configuration to determine if there are any changes to the local BGP routing state (or “Local Routing Information Base”, or “Loc-RIB” to again draw from the conceptual framework). Any changes to the Loc-RIB generate changes in the local state that are passed to BGP neighbours, so these updates are entered in to the Adj-RIBs-Out for subsequent announcement in BGP Update messages to peer ASes. The local forwarding process also processes these changes to the loc-RIB to determine what, if any, changes need to be made to the local forwarding table. (Figure 2)

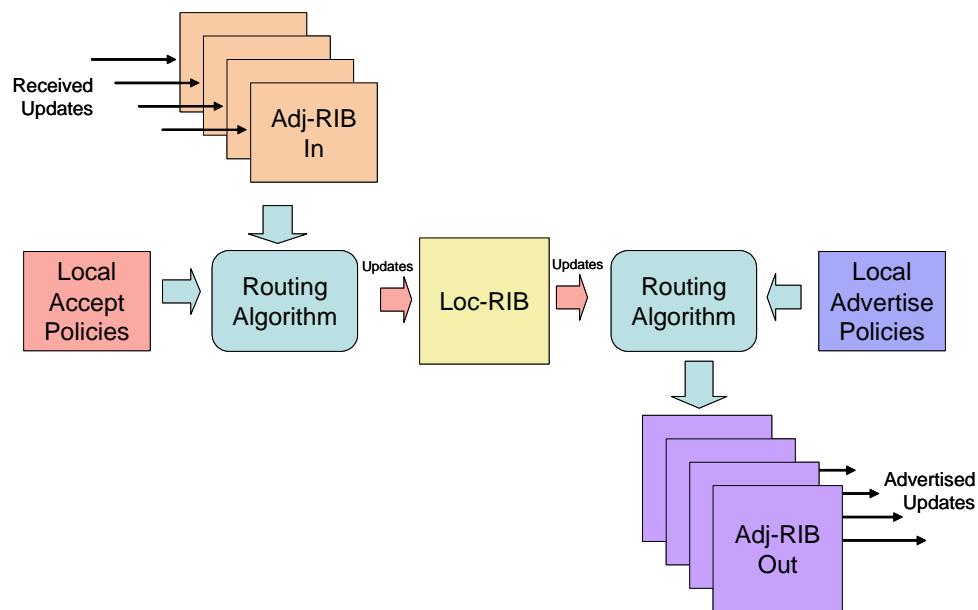


Figure 2 BGP Update Processing

So what are the assumptions that are made in the processing of a BGP Update message? The basic assumption is that the information provided in the update message is correct, and reflects the current and future routing state as seen by the BGP neighbour. The information relating to reachability of address prefixes listed in the Update message will remain valid until the path attribute information is altered in a new Update message, or the address prefix is withdrawn. BGP does not periodically refresh its state, so that once a prefix is described as reachable with a set of AS Path attributes, no further updates will be sent in the lifetime of the BGP session. This implies that the announced prefixes in the update part of the message are valid prefixes reflecting valid address space. It also implies that the Origin AS attribute correctly describes the AS who originated this route, and that the prefix 'owner' has authorized the Origin AS to originate this route. The update implies that the Next Hop address reflects a reachable valid address. Accepting a BGP Update also implies that the AS path reflects a set of implicit permissions, in that each AS pair in the AS path represents a routing permission, where the preceding AS has authorized the next AS to use the preceding AS as transit toward the route origin. Also, the withdrawals from the peer BGP speaker should only reflect previous updates from the same peer, so that the withdrawal should only take effect on any previous announcements from the same neighbour. The problem here is that it is not possible to directly validate these assumptions in a structured manner. How can a transit AS test the implicit assertion that the address prefix controller has allowed the originating AS to generate a routing entry for the block in question, for example.

What can be done to address this level of insecurity in the inter-domain routing space? As a basic wish list of requirements, what would be of value here is to allow a methodology of being able to independently validate the implicit assertions of authority that are associated with a routing advertisement, and to do so in a fashion that could be fully automated in an efficient manner. One possible approach is to have every originator of routing information 'sign' the information in a manner that does not permit third party tampering with the information or the signature. While this does not prevent withholding of information, nor spurious generation of withdrawals and updates, universal adoption of this form of secure attestation of the validity of information could prevent the deliberate injection of false information into the routing system.

Certification and Routing Security

This requirement for tamper-proof signing of information is a well known requirement in secure systems. A common approach to forming such a signature is to generate a *hash*, or a *message digest* of the information. In a secure two-way communication the information can be concatenated with a secret key prior to generating the hash function. The receiver performs a similar function with the secret key and if the message digest values match then the recipient has an increased level of confidence that the message has not been tampered with. But what if you wanted to sign a message such that any recipient would be able to tell that the contents have not been altered?

One possible approach to validate these assertions is to use the properties of public key encryption, where a message block is encrypted with a private key and can only be decrypted by the corresponding public key. The originator of the information could sign the information with their private key, using their private key to encrypt the message digest. The ability to decrypt the

message block with the corresponding public key of the originator, by decrypting the signature block with the public key, and checking that this matches the message digest, would indicate the authenticity of the information, in that that the claimed author really had generated this information, as well as its integrity, in that the information has not been subsequently altered by a third party.

Relating this to the operation of BGP, it would be highly desirable to be able to use an associated validation process to confirm the authenticity of the address prefix described in a BGP Update message before it is accepted into the local RIB. This validation should also extend to validation of the originating Autonomous System, and validation of the implicit binding of the AS to the address prefix as a valid and duly authorized point of origin for the address prefix. By associating digitally signatures with the originator of routing information a third party would not be able to inject false information in to the routing system in that the third party attempting the injection of information would need knowledge of the private key of the purported point of information origination. This looks promising as an approach to the problem, but how can trust relating to the integrity of the association between public keys and routing party identities be maintained?

One possible approach is to use a certificate system with trusted certificate authorities. These authorities need to describe the identity of the party to whom address and AS number resources were allocated, and the public key of the recipient. If the originator of the routing information provides this certificate with the signed routing information, then the recipient can vouch for the integrity of the routing information, provided that the certificate is valid, the identities of the originator of information and the certificate match, the address and originating AS lies within the bounds of the certificate and, lastly, that the and that public key matches the digital signature.

X.509 Certificates

At this point it may be helpful to provide a brief summary of digital certificates.

In any communications network its useful to be able to validate the identity of the other party. One approach to this problem of establishing credentials in a network uses public key cryptography. The essential aspect of this form of cryptography is a pair of encryption keys, normally termed 'public' and 'private' keys. These keys share an interesting property in that a message encrypted with one key can only be decrypted with the other key, and knowledge of one key, with or without an encrypted message, will not provide any information to allow a third party to deduce the value of the other key. So if Alice sends Bob a message encrypted using Alice's private key, and Bob knows Alice's public key, then the fact that Bob can decrypt the message using Alice's public key allows Bob to conclude that the message could only have been sent by Alice. Bob still does not know Alice's private key of course. Indeed Alice could go one step further and sign the message using her private key, followed by an encryption pass using Bob's public key. In this case once Bob decrypts the message using his private key, and then performs a signature check using Alice's public key, then Bob knows that Alice must have sent the message and only Bob can read it. The advantage of this use of public and private keys is that each party keeps their private key secret, then there is no requirement for any form of key exchange as a prerequisite for secure communication. Given that key distribution is perhaps one of the more challenging tasks in any secure system, the use of private / public key pairs is a significant breakthrough for secure communications.

But what if someone wants to impersonate Alice to Bob? Lets think about Carol, who wants to disrupt the communication from Alice to Bob. Carol intercepts the message where Alice attempts to inform Bob of her public key, and Carol informs Bob about her public key instead. In such a case Carol can then masquerade as Alice to Bob, and Bob will be unaware of the switch.

Certificates can fulfill some role here in preventing various forms of identity theft. Certificates introduce a third party, a Certificate Authority (CA), whose role is to issue a certificate that associates the identity of a named party with their public key, and also provides information about the certificate issuer, the time limits that may apply to the use

of the certificate, and the operations for which use of the certificate is valid. Certificates are public attestations by a third party of an association between an identity and a public key, and need not be maintained as a secret. If Bob receives, and trusts, a certificate describing Alice from a Certificate Authority that Bob trusts, then Bob can now store the details of Alice's public key as described in the certificate, and use this to confirm any subsequent communication from Alice that Alice has signed with her private key.

A Certificate Authority (CA) is an organization whose functions is to issue certificates. Its role is to perform a confirmation of the identity of the party who is the subject of the certificate, and attest that the public key in the generated certificate is indeed the public key of the identified party. As well as certification functions, the CA must also provide access to the certificates, while ensuring that the certificates cannot be changed or revoked by any unauthorized party. In addition, a CA provides a capability to revoke a certificate. Revocation can be appropriate in the case of suspicion that a third party has obtained access to the private key, or the private key has been lost, or where the business relationship between the party and the CA has lapsed.

One of the more common forms of certificate, PKIX, uses the X.509 standard. These X.509 certificates can be used for a variety of purposes, as this certificate format (or at least version 3 of this format) allows various forms of extensions to the base certificate structure. X.509 certificates can be used to support public / private key functions for encrypting data, for signing messages, for verifying a signature, as a replacement to a password-based access system for access authentication, and for use within a secure session layer protocol to allow encrypted communication.

Each certificate consists of a certificate body, a signature algorithm and CA signature itself. The certificate body includes a number of common fields and one or more additional fields depending on certificate version. The certificate fields include a version number, a certificate serial number, the starting and ending validity dates of the certificate, the subject of the certificate, the public key of the subject, the CA identification, the CA's signature for the certificate, and one or more certificate extensions.

Alice can present to Bob her certificate. Bob can check the validity of the certificate by checking that the issuer of the certificate is an entity that Bob trusts, or is sufficiently related to one of Bob's trust points that he is prepared to accept the bona fides of the issuing body. Bob can then check that the certificate has not expired, and can check with the certificate issue that the certificate has not subsequently been revoked. At this point Bob can be reasonably sure that he is in possession of Alice's public key, and can use this to validate any communication with Alice.

The basic elements of an X.509 certificate are the issuer, the subject, and the subject's public key. X.509 contains significantly more than this and the full set of certificate information has the following fields:

Version: The valid versions are 1, 2 or 3, encoded, somewhat idiosyncratically, as 0, 1 and 2 respectively

Serial Number: A number, when combined with the Issuer value forms a unique identifier for this certificate

Signature: Not a signature at all, but an identification of the algorithm used to generate the certificate signature, and the value of any associated algorithm parameters. Again one of those idiosyncrasies of the X.509 specification.

Issuer: the name of the certificate issues (an X.500 name, of course!)

Validity: the time when the certificate became valid, and the time when the certificate expires (using a rather verbose ASN.1 notation than manages to encode times between

1950 and 2049 using a 15 octet encoding. At this point its getting harder to define the difference between idiosyncratic and inane!)

Subject: The name of the party whose public key is being certified. This is the X.500 name. It is also possible not to specify this field, and use the extension field of SubjectAltName, which allows DNS names to be used as subject identities.

SubjectPublicKeyInfo: The key algorithm identifier, and the subject's public key,

IssuerUniqueIdentifier: The issuer identification. Deprecated in PKIX

SubjectUniqueIdentifier: A unique subject identification. Deprecated in PKIX

AlgorithmIdentifier: A repetition of the Signature field.(another idiosyncrasy?)

Encrypted: A digital signature that encompasses all but the AlgorithmIdentifier fields.

Extensions: Of the extensions supported by Version 3 of X.509, the extensions of relevance in the routing domain include:

KeyUsage: A bit vector of permitted key usage operations, including the use of the key as a key-signing key, key encipherment (where the key is used to encrypt a message or session key), data encipherment (where the key is used to encrypt data), the Certificate Authority permission, allowing the subject to sign X.509 certificates, and CRL permission, allowing the subject to sign Certification Revocation Lists.

SubjectAltName: A sequence of names that are an alternative to the Subject X.500 name. This allows the use of DNS names in PKIX certificates.

IssuerAltName: An alternate name for the Certificate Issuer.

IPAddrBlock: A list of address prefixes that the subject of the certificate is authorized to use.

ASNumberBlock: A set of Autonomous System numbers that the subject of the certificate is authorized to use.

BGP Security Requirements

If we were looking at integrating certificates into the routing system, it would also be highly desirable to integrate this validation process into the routing protocol, so that some form of certifying information is passed in BGP along with the route objects they refer to with each update, and that certificate validation is an integral part of the BGP route object acceptance into the local RIB and the readvertisement process. While not at quite the same level of desirability, what would also be of some value is to be able to use certifying information to be able to validate the AS path associated with a route object, in that the path represents a valid sequence of AS transits where each downstream AS has explicitly authorized its upstream to be a transit service provider for the address prefix.

If the routing system were to be augmented with such functionality, what should be retained, or improved in any refined inter-domain routing protocol is the performance of the protocol as a near-real time protocol. The time taken for routing information to be propagated across the network, and the time taken for the routing protocol to converge to a stable state should not be lengthened in any substantive way, nor such the additional functionality contribute to an increase in the frequency of BGP updates. The implication is that any additional overheads associated with validation of BGP routing updates should not impose significant additional delays in route object acceptance and readvertisement. It is also tempting to add that this process should not impose additional processing or memory overheads on the processing elements of the routing infrastructure, but such an

objective is perhaps getting beyond what is feasible. There is an obvious tradeoff here, in that the limiting factors of expansion of the size of the routing table relate to the memory requirements of the number of routing table entries, the processing load associated with the acceptance of updates and withdrawals and the network local of the propagation of routing messages. Increasing the both memory and processing requirements for processing routing information has some impact of the feasible number of entries held in the table.

What should also be retained in the inter-domain routing environment is the concept of inter-domain propagation of reachability as a sequence of 'black-box' policy decisions on the part of each AS. The current operational model of the BGP protocol as one where the local policies relating to relative preference when selecting between alternate forwarding paths, and also relating to the choices made in deciding what route objects to readvertise are local policy decisions that are not necessarily published is one that appears to fit in , with provider business models.

This has been the topic of an IETF Working Group, looking at Routing Protocol Security Requirements, and a draft on BGP security requirements is being prepared (draft-ietf-rpsec-bgpsec). The document poses some fundamental questions about the authenticity of routing information that any secure routing framework would need to address. Specifically:

- Is the originating Autonomous system authorized to propagate the prefix we have received?
- Is the AS path, received via an UPDATE, valid?

The verification of AS-Path validity falls into three distinct categories.

- Does the AS-Path specified actually exist and, based on the AS-PATH, is it possible to traverse that path to reach a given prefix?
- Has the update actually travelled the path?
- Was the update authorized to traverse the given path by the originator of the prefix?

[draft-ietf-rpsec-bgpsec]

The document notes that the requirements for a secure routing system should include the property that its time to reach a stable state, or convergence time, should be unaltered by the inclusion of security functionality, and the values of various state timers should not need to be altered. The requirements document advocates a preference for real time authentication of the contents of Update messages, but also allows for an approach using periodic sweep of the loc-RIB to validate the originating AS and AS-PATH for each prefix in the RIB, and notes that at the least an implementation of a secure system should support one or the other of these approaches.

It is also noted that the time for simultaneous activation of Internet features is long gone, and the only model for deployment of a secure inter-domain routing protocol will be by piecemeal incremental deployment. This implies that the routing system will contain both verifiable and unverifiable information , and allow for systems that support verifiable message exchange to be able to communicate this capability to its peers in a backward-compatible fashion. It is an interesting issue whether a local operator should express a preference for verifiable information in such a heterogeneous environment. Some operators would see this as being a selection preference only when the conventional route discriminators (prefix length, local preference, AS Path length) are equal, while other operators may see this validation capability as being an input to the local preference function. In general, this can be safely left as a local decision, as the potential for loop formation through differing selection policies is eliminated through the protocol's use of explicit AS Path information as a route attribute.

The requirement document notes the potential to use either a strict hierarchical trust model, or a distributed mutual trust model. The author has a strong personal preference in a hierarchical trust model, where the origin of the certificate chain is an allocation certificate describing the Regional Internet Registry, and the IP address and AS number resources that the RIR administers. The RIR is the able to issue certificates that describe the allocation of resources to an ISP or Local Internet Registry, who may use these certificates as validation for routing entries, or, in turn issue certificates describing resource allocations at the next level of the hierarchy, and so on. In this way a certificate bounds the resources that can be described in the certificate such that they must lie within the bounds of the resources set described in the next level up in the hierarchy.

The authentication of update messages involves a number of verification tests. The originating AS must be verified as being an authorized announcer of the address prefix. This requires some form of attestation by the address holder that they are the legitimate unique holder of the address space, and that they have authorized the holder of a particular AS to use that AS to originate a BGP advertisement. The AS Path list must correspond to a verifiable sequence of autonomous systems that correspond to a viable path across the network, where the first element of the AS Path corresponds to the local peer AS. The objective that security information should propagate at the same speed as routing information implies that there is a need to distribute information relating to the validity of AS numbers and address prefixes, information relating to the authorization of an Autonomous System to originate an address announcement, and information relating to AS connectivity, at a rate that is at least as fast as routing update propagation.

The system should provide a method to allow the receiver of a BGP Update to verify that the listed originating AS actually originated the update, and that the AS's listed in the AS_PATH actually forwarded the update in the order described by the AS Path.

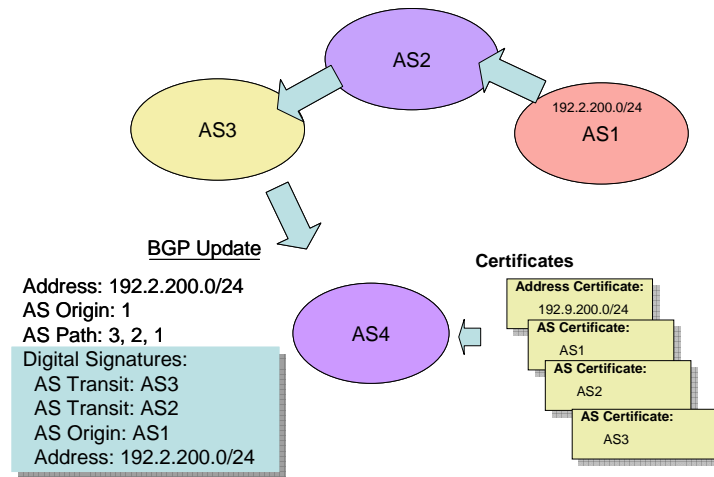


Figure 3 – Processing a BGP Update Message

One form of example of the required processing steps is illustrated in Figure 3. Here AS1 originates an advertisement for the address prefix 192.9.200.0/24. In a secure routing environment AS1 would associate a digitally signed authorization from the holder of this address block that permits AS1 to originate a routing advertisement for this address prefix. In passing this advertisement to AS2 AS1 would sign this advertisement with its own private key, indicating that it is the authorized holder of AS1 and that it is the party that has been authorized to make this advertisement. As the update propagates across the network each transit AS would add its own digital signature upon the signature block, indicating that it has received the update from its AS peer and that it Validation of the Update as received by AS4 would require AS4 to verify the signatures for each of the AS objects in the AS Path to validate that the signing entity is indeed the authorized holder of that AS object., and then to check the signature associated with the Address object to validate the authorization attestation as well as the validity of the address.

The Role of the Regional Internet Registries

Irrespective of how these requirements are implemented in a routing protocol there remains the requirement to provide trust anchors for address and AS information in the form of a Public Key Infrastructure (PKI). Here there is a distinct role for the Regional Internet Registries to augment their established role of address allocation with a complementary role of being a certificate authority for the entities that are the recipients of such allocations. What is being asked of the RIRs in such a scenario is to provide a public record of their allocation actions in the form of a public key certificate that records the allocation of a specific resource to the entity that possesses the matching public/private key pair.

The purpose of the certificate that would be issued by an RIR in this context is subtly different from a more conventional certification of an entity's public key. In this case the information provided by the RIR-issued certificate is that the subject of the certificate, whose public key is included in the certificate, has been allocated the address and AS resources as described in the certificate's extension. To continue the certificate hierarchy, this certificate would include a Key Usage permission to use the key as a key signing key, and for the subject to be permitted to use the key to be a Certificate Authority to be able to issue further certificates. This would allow LIR or ISP to issue further certificates to match further sub-assignments that may be undertaken from the original RIR-assigned address or AS number block. Again, such certificates are not a confirmation of an identity assertion, but are assertions that the holder of the corresponding private key has been duly assigned the right to use the described address and/or AS number block.

This certificate structure would allow the subject of an address allocation to digitally sign an authority to allow a nominated AS to originate an advertisement. The associated certificate allows a third party to authenticate the validity of this authority by virtue of the match of the digital signature in the authority and the public key of the associated certificate, coupled with the confirmation that the address prefix is contained within the bounds of the address extension of the certificate, and the current time is within the validity time of the certificate and the certificate has not been already revoked.

The same certificate structure would allow an entity to attach a digital signature to a BGP update asserting that it has the right to use that AS number as a route object. The associated certificate that includes the AS number as part of the certificate extension, with its matching public key, would allow a third party to validate the AS number as being a valid part of a route object.

Next Steps

There are a number of potential approaches to inclusion of this keying information into the inter-domain routing framework. One of the more long-standing approaches is that of Secure-BGP, a protocol whose development has been largely lead by Steven Kent of BBN Technologies. This is a relatively comprehensive approach to securing routing information that attempts to secure both the injection of address reachability information into the inter-domain routing system, and to secure the integrity of the AS path associated with the address updates. A variation of this approach has been developed by Cisco Systems, Secure Origin BGP, where the major functional difference inn this approach is weakening the verification of the AS Path to a level of "plausibility". This approach attempts to reduce some of the memory and processing overheads associated with AS Path validation by using an approach of generating a validated connectivity graph of AS connectivity, and then checking that each presented AS Path matches a sequence that can be derived from this connectivity graph. Other approaches have advocated out-of-band communications as a means of verification of routing information, and, possibly one of the earliest proposals in this space, the use of the DNS as a means of storing origin validation information.

The IETF appears to be embarking on a procedure that entails the completion of the security requirements activity as a precondition to undertaking any standardization activity relating to a secure inter-domain routing protocol. Whether the IETF can undertake this activity within a timescale that matches the increasing levels of concern about routing security in the Internet remains an open issue.

Of course the actions involving the establishment of trust anchors in a PKI framework for address and AS allocation information need not await the outcome of this IETF activity, and there is an increasingly compelling case to embark on this certification activity now, rather than awaiting further protocol standardization activity. Even without a particular protocol, such certificates can be used within the current process of validation of routing requests, where a client could use their private key to sign a routing request, and attach their associated certificate to the request. The upstream transit service provider can validate the certificate, and, in turn validate the authenticity of the routing request. It may appear to be a small step, but it provides a useable and, importantly, a cost effective process that addresses one of the most obvious vulnerabilities in today's routing environment, that of validation of routing requests in an efficient and timely manner.

Geoff Huston

Disclaimer

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

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 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 is currently the Senior Internet Researcher at APNIC, the Regional Internet Registry serving the Asia Pacific region.