The ISP Column
A monthly column on things Internet

July 2011
Geoff Huston
Randy Bush

Securing BGP with BGPsec

Introduction

For many years the Internet's fundamental elements – names and addresses – were the source of basic structural vulnerabilities in the network. With the increasing momentum behind the deployment of DNSSEC there is some cause for optimism that we have the elements of securing the name space now in hand, but what about addresses and routing? In this article we will look at current efforts within the IETF to secure the use of addresses within the routing infrastructure of the Internet, and the status of current work of the Secure Inter-Domain Routing (SIDR) working group.

We will look at the approach taken by the SIDR Working Group, and examine the architecture and mechanisms that have been adopted as part of this study. This work was undertaken in three stages: the first concentrated on the mechanisms to support attestations relating to addresses and their use; the second looked at how to secure origination of routing announcements; and the third looked at how to secure the transitive part of BGP route propagation.

Supporting Attestations about Addresses through the Resource Public Key Infrastructure (RPKI)

Prior work in the area of securing the Internet's routing system has focused on the operation of the Border Gateway Protocol (BGP) in an effort to secure the operation of the protocol, and validate, as far as is possible, the contents of BGP Update messages. Some notable contributions in more than a decade of study include Secure-BGP (S-BGP) [sBGP], Secure Origin BGP (soBGP) [soBGP], Pretty Secure BGP (psBGP) [psBGP], IRR [IRR], and the use of an AS RR in the DNS, signed by DNSSEC [DNS].

The common factor in this prior work was that they all required, as a primary input, a means of validating basic assertions relating to origination of a route into the inter-domain routing system: that the IP address block and the AS numbers being used are valid and that the parties using these IP addresses and AS numbers in the context of routing advertisement are properly authorized to so do.

The approach adopted by SIDR for the way in which trust is formalized in the routing environment is through the use of Resource Certificates. These certificates are X.509 certificates that conform to the PKIX profile [PKIX]. They also contain an extension field that lists a collection of IP resources (IPv4 addresses, IPv6 addresses and AS Numbers) [RFC3779]. These certificates attest that the certificate's issuer has granted to the certificate's subject a unique "right-of-use" for the associated set of IP resources, by virtue of a resource allocation action. This concept mirrors the resource allocation framework of the IANA (Internet Assigned Numbers Authority), the regional Internet registries (RIRs), operators and others, and the certificate provides a means for a third-party (relying party) to formally validate assertions related to resource allocations [sidr-arch].

The hierarchy of the RPKI is based on the administrative resource allocation hierarchy, where resources are distributed from the IANA to the RIRs, to Local Internet Registries (LIRs) and end users. The RPKI mirrors this allocation hierarchy with certificates that match current resource allocations (Figure 1).

The Certification Authorities (CAs) in this RPKI correspond to entities that have been allocated resources. Those entities are able to sign authorities and attestations, and to do so they use specific purpose End Entity (EE) certificates. This additional level of indirection allows the entity to customize each issued authority for specific subsets of number resources that are administered by this entity. Through the use of single-use EE certificates, the issuer can control the validity of the signed authority through the ability to revoke the EE certificate used to sign the authority. As is often the case, a level of indirection comes in handy.

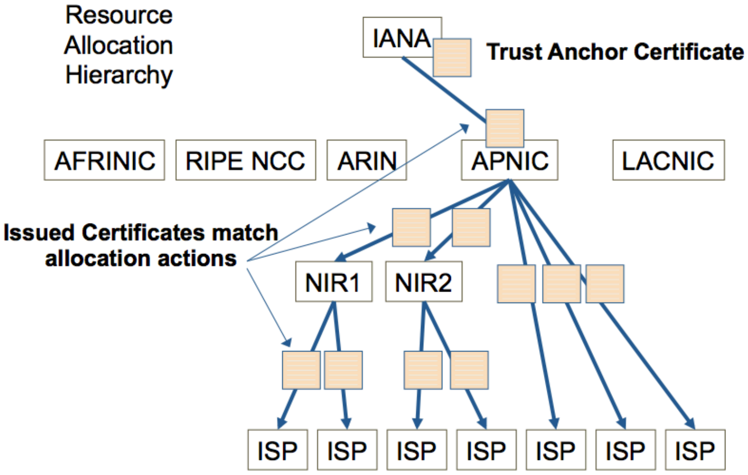

Figure 1 – Hierarchy of the RPKI

Signed attestations relating to addresses and their use in routing are generated by selecting a subset of resources that will be the subject of the attestation, by generating an EE certificate that lists these resources, and by specifying validity dates in the EE certificate that correspond to the validity dates of the authority. The authority is published in the entity's RPKI repository publication point. The RPKI makes conventional use of Certificate Revocation Lists (CRLs) to revoke certificates that have not expired, but which are no longer valid. Every CA in the RPKI regularly issues a CRL according to the CA's declared CRL update cycle. A CA certificate may be revoked by an issuing authority for a number of reasons, including key rollover, the reduction in the resource set associated with the certificate's subject, or termination of the resource allocation. To invalidate an object that can be verified by a given EE certificate, the CA that issued the EE certificate can revoke the corresponding EE certificate.

The RPKI uses a distributed publication framework, wherein each CA publishes its products (including EE certificates, CRLs and signed objects) at a location of its choosing. The set of all such repositories form a complete information space, and it is fundamental to the model of securing BGP in the public Internet that the entire RPKI information space is available to every Relying Party (RP). It is the role of each RP to maintain a local cache of the entire distributed repository collection by regularly synchronizing each element in the local cache against the original repository publication point. To assist RPs in the synchronization task, the each RPKI publication point uses a manifest. A manifest is a signed object that lists the names (and hash values) of all the objects published at that publication point. It is used to assist RPs to ensure that they have managed to synchronize against a complete copy of the material published at the CA's publication point.

The utility of the RPKI lies in its ability to validate digitally signed information and, therefore, give relying parties some confidence in the validity of signed attestations about addresses and their use. The particular utility of the RPKI is not as means of validation of attestations of an individual's identity or their role, but as a means of validating their authority to use IP address resources. While it is possible to digitally sign any digital object, it has been suggested that the RPKI system uses a very small number of standard signed objects that have particular meaning in the context of routing security.

Securing Route Origination

The approach adopted by SIDR to secure origination of routing information is one that uses a particular signed authority, a Route Origination Authorization (ROA) [ROA]. An ROA is an authority created by a prefix holder that authorizes an AS to originate one or more specific route advertisements into the inter-domain routing system. An ROA is a digital object formatted according to the Cryptographic Message Syntax specification (CMS) [RFC3852] that contains a list of address prefixes and one AS number. The AS is the specific AS being authorized to originate route advertisements for one or more of the address prefixes in the ROA. The CMS object also includes the EE resource certificate for the key used to verify the ROA. The IP Address extension in this EE certificate must encompass the IP address prefixes listed in the ROA's contents.

The ROA conveys a simple authority. It does not convey any further routing policy information, nor does it convey whether or not the AS holder has even consented to actually announce the prefix(es) into the routing system. The associated EE certificate is used to control the validity of the ROA and the CMS wrapper is used to bind the ROA and the EE certificate within a single signed structure in a secure fashion.

There is one special ROA, one that authorizes AS 0 to originate a route. As AS 0 is a reserved AS that should never be used by a BGP speaker, this ROA is a "negative" authority, used to indicate that no AS has authority to originate a route for the address prefix(es) listed in the ROA.

If the entire routing system were to be populated with ROA's, then identification of an invalid route advertisement would be directly related to detection of an invalid ROA or a missing ROA. However in a more likely scenario of partial use of ROA's (such as when only some legitimate route originations are authorized in a ROA), the absence of an ROA cannot be interpreted simply as an unauthorized use of an address prefix. This leads to the use of a tri-state validation process for routes. If a given route matches exactly the information contained in an ROA whose EE certificate can be validated in the RPKI (a "valid" ROA) then the route can be regarded as a "valid" origination. Where the address prefix matches that in a valid ROA, but the origination AS does not match the AS number in the ROA, and there are no other valid ROAs that explicitly validate the announcing AS, then the route can be considered to be "invalid". Also, where the address prefix is more specific than that of a valid ROA, and there are no other valid ROAs that match the prefix, then the route can also be considered "invalid". Where the prefix in a route is not described in any ROA and is not a more specific prefix of any ROA, then the route has an "unknown" validation outcome. These three potential outcomes can be considered a set of relative local preferences. Routes whose origin can be considered "valid" are generally proposed to be preferred over routes that are unknown, which, in turn, can generally be preferred over routes that are considered invalid. However, such relative preferences are a matter to be determined by local routing policy. Local policies may choose to adopt a stricter policy and, for example, discard routes with an invalid validation outcome [sidr-roa-validation].

The way in which ROAs are used to validate the origin of routes in BGP differs from many previous proposals for securing BGP. In this framework the ROAs are published in the RPKI distributed repository framework. Each RP can use the locally cached collection of valid ROAs to create a validation filter collection, with each element of the set containing an Address, prefix size constraints and an originating AS. It is this filter set, rather than the ROAs themselves that are fed to the local routers [sidr-rpki-router]. (An example of the way in which ROAs can be used to detect prefix hijack attempts is shown in Figure 2)

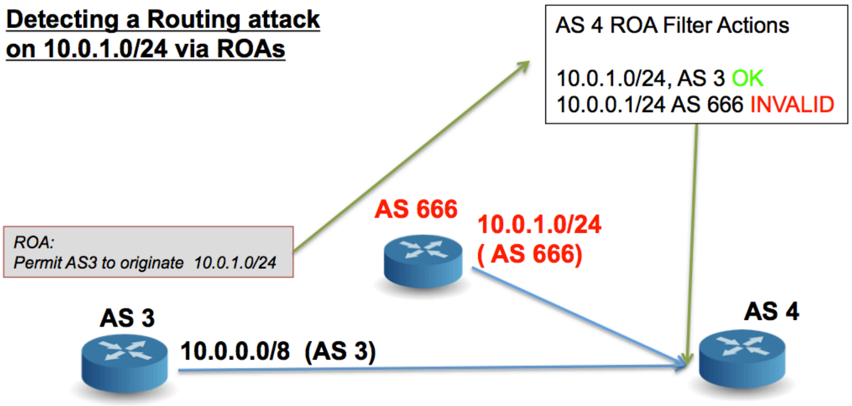

Figure 2 – Use of ROAs to detect unauthorised Route Origination

The model of injecting validation of origination into the BGP domain is an example of a highly modular and piecemeal deployment. There are no changes to the BGP protocol for this origin validation part of the secure routing framework.

The process of securing origination starts with the address holder, who generates local keys and requests certification of their address space from the entity from whom their addresses were allocated or assigned. With this CA resource certificate, the address holder is then in a position to generate an EE certificate and a ROA that assigns an authority for a nominated AS to advertise a route for an address prefix drawn from their address holdings. The one condition here is that if an address holder issues a ROA for an address prefix providing an authority for one AS to originate a route for this prefix, then the address holder is required to issue ROAs for all the AS's that have been similarly authorized to originate a route for this address prefix. The address holder publishes this ROA in their publication point in the distributed RPKI repository structure.

Relying Parties can configure a locally managed cache of the distributed RPKI repository and collect the set of valid ROAs [rcynic]. They can then, via the dedicated RPKI cache-to-router protocol [rpki-rtr], maintain, on a set of "client" routers the set of address prefix/originating AS authorities that are described in valid ROAs. This information can be used by the BGP-speaking router as an input to the local route decision process.

This model of operation supports piecemeal incremental deployment, wherein individual address holders may issue ROAs to authorized routing advertisements independent of the actions of other address holders. Also, ASs may deploy local validation of route origination independently of the actions of other ASs. And given that there are no changes to the operation of BGP, then there are no complex interdependencies that hinder piecemeal incremental deployment of this particular aspect of securing routing.

Securing Route Propagation - BGPsec

Origin validation as described earlier does not provide cryptographic assurance that the origin AS in a received BGP route was indeed the originating AS of this route. A malicious BGP speaker can synthesize a route as if it came from the authorized AS. Thus, it is very useful in detecting accidental misannouncements, but origination validation does little to prevent malicious routing attacks from a determined attacker.

In looking at the operation of the BGP protocol, some parts of the protocol interaction are strictly local between two BGP-speaking peers, such as advising a peer of local attributes. Another part of the BGP protocol is a "chained" interaction, in which each AS adds information to the protocol object. This attribute of a BGP update, the AS Path, is not only useful to detect and prevent routing loops, it is also used in the BGP best path selection algorithm.

A related routing security question concerns the validity of this "chained" information, namely the AS Path information contained in a route. Within the operation of the BGP protocol, each AS that propagates an update to its AS neighbours is required to add its AS number to the AS Path sequence. The inference is that at any stage in the propagation of a route through the inter-domain routing system, the AS Path represents a viable AS transit sequence from the local AS to the AS originating the route. This AS Path attribute of a route is used for loop detection. Locally, the AS Path may also be used as an input to a local route policy process, using the length of the AS Path as route metric.

Attacks on the AS path can be used to subvert the routing environment. A malicious BGP speaker may manipulate the AS Path to prevent an AS from accepting a route by adding its AS number to the AS Path, or it may attempt to make a particular route more likely to be selected by a remote AS by stripping out AS's from the AS Path. Accordingly, it is important to equip a secure BGP framework with the ability to validate the authenticity of the AS Path presented in a BGP update [kapela/pilosov].

In attempting to validate an AS path there are a number of potential validation questions.

- The first and weakest question is: Are all ASs in the AS Path valid ASs?

- A slightly stronger validation question is: Do all the AS pairs in the AS Path represent valid AS adjacencies (where both ASs in the pair-wise association are willing to attest to their mutual adjacency in BGP)?

- A even stronger question is: Do the sequence of ASs in the AS Path represent the actual propagation path of the BGP route object?

This last question forms the basis for the SIDR activity in defining an AS Path validation framework, BGPsec. This is an attempt to assure a BGP speaker that the operation of the BGP protocol is operating correctly and that the content of a BGP update correctly represents the inter-AS propagation path of the update from the point of origination to the receiver of the route. This is not the same as a policy validation tool and it does not necessarily assure the receiver of the route that this update conforms to the routing policies of neighbouring BGP speakers. This route also does not necessarily reflect the policy intent of the originator of the route. The BGPsec framework proposed for securing the AS Path also makes use of a local RPKI cache, but it includes an additional element of certification. The additional element of the security credentials used here is an extension to the certification of AS numbers with a set of operational keys and their associated certificates used for signing update messages on eBGP routers in the AS. These "router certificates" can sign BGP update attributes in the routing infrastructure, and the signature can be interpreted as being a signature made "in the name of" an AS number.

In the BGPsec framework, eBGP speaking routers within the AS have the ability to "sign" a BGP update before sending it. In this case, the added signature "covers" the signature of the received BGP update, the local AS number, the AS number to which the update is being sent, as well as a hash of the public key part of the router's key pair used to sign route updates. The couplet of the public key hash and the signature itself is added to the BGP protocol update as a BGPsec update attribute. As the update traverses a sequence of transit ASes each eBGP speaker at the egress of each AS adds its own public key hash and digital signature to the BGPsec attribute sequence (Figure 3).

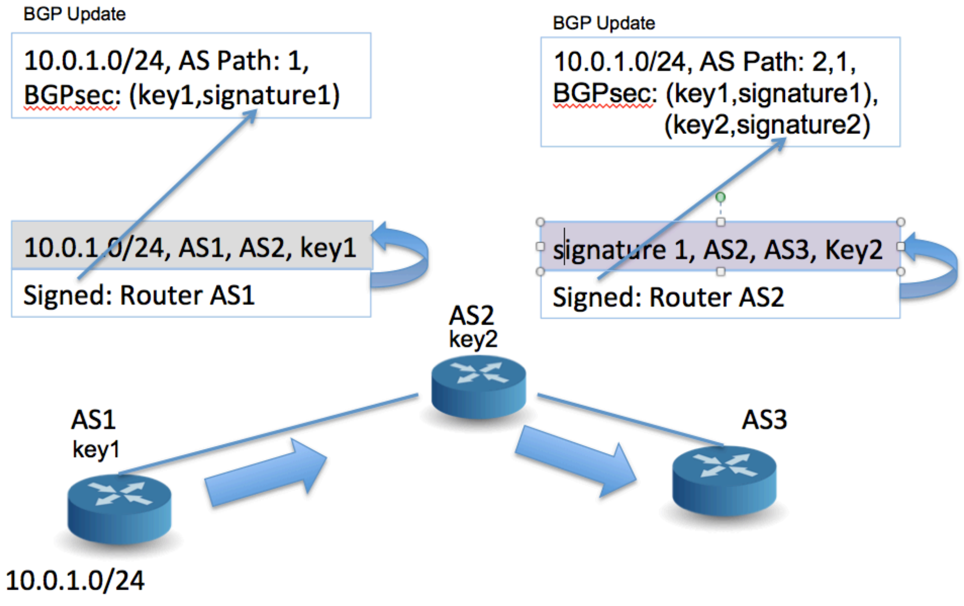
BGP Update
10.0.1.0/24, AS Path: 1,
BGPsec: (key1,signature1)
BGP Update
10.0.1.0/24, AS Path: 2,1,
BGPsec: (key1,signature1),
(key2,signature2)
10.0.1.0/24, AS1, AS2, key1
Signed: Router AS1
signature 1, AS2, AS3, Key2
Signed: Router AS2
AS2
key2
AS1
key1
AS3
10.0.1.0/24

Figure 3 – BGPsec AS Path Protection

This interlocking of signatures allows a receiver of a BGP update to use the interlocking chain of digital signatures to validate (for each AS in the AS Path) that the corresponding signature was correctly generated "in the name of" that AS in the AS path, and that the next AS in the path matches the next AS in the signed material. The "forward signing" that includes the AS to which the update is being sent prevents a man-in-the-middle attack of the form of taking a legitimate outbound route announcement destined for one neighbour AS and redirecting it to another AS. But this signing of the AS Path is not quite enough to secure the route update, as the AS Path needs to be coupled to the actual address prefix by the originator of the route. The route originator needs to sign across not only the local AS and the AS to whom the route update is being sent, but also the address prefix and the expiry time of the route. This allows the path to be "bound" to the prefix and prevents a man-in-the-middle splicing a signed path or signed path fragment against a different prefix.

If the signatures that "span" the AS Path in the BGP update can all be validated, then the receiver of the BGP update can validate, in a cryptographic sense, the currency of the routing update. It can also validate that the route update was propagated across the inter-AS routing space in a manner that is faithfully represented in the AS Path of the route.

The expiry time of the EE certificates used in conjunction of signed route updates introduces a new behaviour into BGPsec. In the context of BGP, an announced route remains current until it is explicitly withdrawn or until the peer session that announced the route goes down. This property of BGP introduces the possibility of "ghost route" attacks in BGP, wherein a BGP speaker fails to propagate a withdrawal in order to divert the consequent misdirected traffic from its peers. In BGPsec, all route advertisements are given an expiry time by the originator of the route. This expiry time corresponds to the notAfter time of the EE certificate used to sign the protocol update, after which time the route is to be considered invalid. The implication is that an originator of a route is required to re-advertise the route, and refresh the implicit expiry timer of the associated digital signature at regular intervals.

This approach to route update validation is not quite the "light-touch" of origination validation. In this case the mechanism requires the use of a new BGP attribute and negotiation of a new BGP capability between eBGP peers. In turn, this means that the model of incremental deployment is one that is more "viral" than truly piecemeal. By "viral" we mean that this is a model of incremental deployment in which direct eBGP peers of a BGPsec-speaking AS will be able speak BGPsec between themselves in a meaningful way. In turn these adjacent AS's can offer to speak BGPsec with their eBGP peers, and so on. This does not imply that BGPsec deployment must necessarily start from a single AS, but it does imply that communities of interconnected AS's all speaking BGPsec will be able to provide assurance via BGPsec on those routes originated and propagated within that community of interconnected ASs. It also implies that the greatest level of benefit to adopters of secure BGP will be realized by ASs that adopt BGPsec as a connected community of ASs.

There are other changes to the behaviour of BGP that are implied by this mechanism. BGP conventionally permits "update packing" where a number of address prefixes can be placed in a single update message if they share a common collection of attributes, including the AS Path. At this stage it appears that such update packing would not be supported in secure BGP, and each update in secure BGP would refer to a single prefix. Obviously this would have some impact on the level of BGP traffic, but early experiments suggest not at an unreasonable cost.

There are further impacts on BGP that have not been fully quantified in studies to date. The addition of a compound attribute of a signature and a public key identifier for every AS in the AS Path has size implications on the amount of local storage a secure BGP speaker will need to store these additional per-prefix per-peer attributes. It also has broader implications if used in conjunction with current proposals for multi-path BGP where multiple paths, in addition to the "best" path are propagated to eBGP peers. Also, the computational load of validation of signatures in secure BGP is significantly higher in terms of the number of cryptographic operations that are required to validate a BGP update.

However, BGPsec is not intended to "tunnel" across those parts of the inter-domain routing space that do not support BGPsec capabilities. When an update leaves a BGPsec realm, the BGPsec signature attributes of the route are stripped out, so the storage overheads of BGPsec are not seen by other BGP speakers. Similarly, the periodic updates that result from the expiry timer should not propagate beyond the BGPsec realm. If the boundary is prepared to perform BGP update packing to non-BGPsec peers then even the unpacked update overhead is not carried outside of the BGPsec realm.

It is also noted that the "full" load of BGPsec would only necessarily be carried by "transit" ASs; that is, those ASs that propagate routes on behalf of other ASs. Historically we see some 15 percent of ASs are "transit" ASes, while all other ASes behave as "stub" ASes that only originate routes and do not appear to transit routes for others. Such stub ASes can support a "light weight" simplex version of BGPsec that can either point default a default route to its upstream AS provider, or trust its upstream ASs to perform BGPsec validation. In this case the stub AS needs to provide BGPsec signed originated routes to its upstream ASs, but no more.

Conclusion

The work on the specification of the RPKI itself and the specification of origin validation is nearing a point of logical completion of the first phase of standardization within the IETF, and the working draft documents are being passed from the working group into the review process leading to their publication as proposed standard RFCs. The RIRs are in the process of launching their RPKI services based on these specifications and the initial deployment of working code has been made by a number of parties, who are also working on integration of origination validation in BGP implementations.

The work on securing the AS Path is at an earlier phase in the development process and the initial design material is being considered by the SIDR Working Group. It is expected to take a similar path of further review and refinement in light of developing experience and study of the proposed approach.

The RPKI has been designed as a robust and simple framework. As far as possible, existing standards, technologies, and processes have been exploited, reflecting the conservatism of the routing community and the difficulty in securing rapid, widespread adoption of novel technologies.

Acknowledgements

The work described here is the outcome of the efforts of many individuals who have contributed to securing BGP over a period that now spans two decades, and certainly too many to ensure that all the contributors are recognized here. Instead, the authors would like to acknowledge their work and trust that the mechanisms described here are a faithful representation of the cumulative sum of their various contributions.

References

[sBGP] S. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol (S-BGP)," IEEE Journal on Selected Areas in Communications, vol. 18, no. 4, pp 582-592, April 2000.

[soBGP] R. White, "Securing BGP through secure origin BGP," Internet Protocol Journal, vol. 6, no. 3, September 2003.

[psBGP] P. van Oorschot, T. Wan and E. Kranakis, "On Interdomain Routing Security and Pretty Secure BGP (psBGP)," ACM Transactions on Information and System Security, vol. 10, no. 3, July 2007.

[IRR] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis and P. McDaniel, "Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing," Proc. of Internet Society Symposium on Network and Distributed System Security (NDSS‚Äô03), February 2003.

[DNS] T. Bates, R. Bush, T. Li and Y. Rekhter, "DNS-based NLRI origin AS verification in BGP," Internet Draft, July 1998.

[PKIX] D. Cooper et al., "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile," Request for Comment RFC5280, May 2008.

[RFC3779] C. Lynn, S. Kent and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers," Request for Comment RFC3779, June 2004.

[sidr-arch] M. Lepinski, S. Kent, "An Infrastructure to Support Secure Internet Routing," work in progress (Internet Draft), February 2008.

[sidr-cert-profile] G. Huston, G. Michaelson, R. Loomans, "A Profile for X.509 PKIX Resource Certificates," work in progress (Internet Draft), September 2008.

[ROA] M. Lepinski, S. Kent, D. Kong, "A Profile for Route Origin Authorizations (ROAs)," work in progress (Internet Draft), July 2008.

[RFC3852] R. Housley, "Cryptographic Message Syntax (CMS)," Request for Comment RFC3852, July 2004.

[sidr-rpki-router] R. Bush, R. Austein, "The RPKI/Router Protocol", work in progress (Internet Draft), March 2011.

[kapela/pilosov] http://www.wired.com/threatlevel/2008/08/revealed-the-in/, August 2008.

Disclaimer

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