Resource Certification - A Public Key Infrastructure for IP Addresses and AS's

Geoff Huston, George Michaelson and Stephen Kent

Resource Certification - A Public Key Infrastructure for IP Addresses and AS's

Geoff Huston, George Michaelson and Stephen Kent

Abstract— **We examine a form of an X.509 Public Key certificate that is used to bind IP address and AS number resources to a public/private key pair. These certificates are used to attest to resource allocation actions, so that digitally signed attestations relating to a party's right-of-use of IP addresses and AS numbers can be validated by relying parties, using a related Resource Certificate Public Key Infrastructure. This has particular application in the area of demonstrable attestations related to the right-of-use of IP addresses, and in the area of inter-domain routing security. The issues related to the application of this PKI to inter-domain routing security are considered, and the design, management and use of resource certificates, and the structure of the related Public Key Infrastructure are described in detail.**

Index Terms— **BGP Security, Inter-Domain Routing Security, Public Key Infrastructure, X.509**

I. Introduction

IN November 2008 the Asia Pacific Network Information Centre (APNIC) announced the release of a public resource certification service that makes use of X.509 technology [1] to publish public key certificates and associated signed objects that uniquely associate a private key holder with a 'right-of-use' of a collection of IP number resources (IPv4 addresses, IPv6 addresses and Autonomous System (AS) Numbers). This APNIC activity forms part of a larger certificate infrastructure effort that is ultimately intended to provide certification for all number resources in the public Internet. This report describes this Resource Public Key Infrastructure (RPKI) in more detail, looking at the various aspects of the design that lie behind the construction of this particular PKI.

The objective of the RPKI is to provide a means of validating the authenticity of certain types of assertions about an IP address or AS. This authenticity means being able to determine that an address or AS number has been validly allocated or assigned, that the address can be announced into the Internet's inter-domain routing system, and that the AS number can be used within the attributes of the routing information system's object set. In addition, the RPKI can validate the association between an address or AS number and its current right-of-use holder.

Manuscript received July 1, 2009. This work was supported by the Asia Pacific Network Information Centre (APNIC) and BBN *Technologies*.

G. Huston is with APNIC, Milton, QLD 4064 Australia (phone: +61-7-38583100; e-mail: gih@apnic.net).

G. Michaelson is with APNIC, Milton, QLD 4064 Australia (e-mail: ggm@apnic.net).

S. Kent is with BBN Technologies, Cambridge, MA, USA, (e-mail: kent@bbn.com)

II. Prior Work in Routing Security

The initial approach used to provide some level of ability to determine the legitimacy of the use of IP addresses in the routing system was the IP address allocation registry, administered by the five Regional Internet Registries (RIRs). The RIRs' registries collectively contain the current list of all validly allocated number resources and the details of the identity of the party to whom the resources were allocated.

There are some problems in using this published registry information, in that the registry data is not published in a complete format, it is incomplete and inconsistent in places, and the query tool, "whois" [2], is insecure and readily disrupted by a number of forms of attack.

This registry approach was refined in the development of "Internet Routing Registries" (IRRs). An IRR database contains entries that relate to the inter-AS adjacencies that exist in the routing space, and the applicable routing policies that apply to these adjacencies. It also contain entries that describe origination of routing information, binding an address prefix to an originating AS. IRRs use the RPSL [3] notation to describe routing policies. The major operational use for IRRs has been in the automated construction and maintenance of routing filters for routers operating at the boundaries between ASes. (These routers implement the Border Gateway Protocol (BGP) [4] and thus they are often referred to as "border routers"). By processing the data in an IRR, matching AS import and export routing policies and joining the inferred propagation information to the IRR-declared prefix origination for each AS, it is possible to construct the list of all prefixes that an adjacent AS may announce to its peer. From that information, a local filter can be constructed, that allows the local BGP instance the ability to declare any other routing information as unauthorized and filter it out of consideration.

The IRR framework is intended to ensure that routes are added into the routing system via a deliberative operational process, rather than as a potentially accidental or malicious outcome. However, IRRs are not used universally, and the partial use of IRR systems limits their general applicability. This approach has experienced a number of problems, including the inability to authenticate the data retrieved from an IRR, the number of IRRs and the diversity of policies of data admission and the inconsistencies between each IRR. The IRR publication model is not inherently secure and very few

IRRs implement a strict condition that IRR data should be derived from allocation registry data. There is no easy method for a client of an IRR to establish the currency and accuracy of IRR data [5].

The trust model of the IRRs appears to relate to trust in the data admission policies of the IRR, which, in turn, places an undue level of reliance in the location of publication of the data, as distinct from establishing trust through explicit validation of the data. Efforts to improve this situation were studied in the late 1990s, but few IRRs have implemented the measures proposed by this Routing Policy System Security study [6].

Prior work has also focused on the operation of BGP in an effort to secure the operation of the protocol and validate the contents of BGP Update messages. Some major contributions in this area of study so far include S-BGP [7], soBGP [8], psBGP [9] , IRR [10], and the use of an AS RR in the DNS, signed by DNSSEC [11].

The common factor in these approaches is that they all require, as a primary input, a means of validating two basic assertions relating to origination of a route into the inter-domain routing system: firstly, that the IP address block and the AS numbers being used are valid, and, secondly, that the parties using these IP addresses and AS numbers are properly authorized to so do.

The mechanisms proposed to perform this validation vary from simple assertion through peer corroboration to use of a comprehensive resource PKI. Those proposals that rely on the existence of a comprehensive resource PKI do so despite the obvious fact that no such authoritative and comprehensive PKI exists today. Where the proposals make use of weaker models of assertion and/or a web of trust, such mechanisms could be replaced by a resource PKI with no loss of functionality, and with a significant improvement in the level of trust that could be placed in the outcome of the validation process. The essential common approach across all of these proposals to secure BGP is to provide a "feed" of signed credential information, which can be used to validate the feed of routing information, where the validation of these credentials could be undertaken by a resource PKI.

III. Resource Certificates and the Resource Public Key Infrastructure

Resource Certificates are X.509 certificates that conform to the PKIX profile [12] and that also contain a mandatory certificate extension that lists a collection of IP resources (IPv4 addresses, IPv6 addresses and AS Numbers) [13]. These certificates attest that the certificate's issuer has granted to the 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, where the certificate provides a means of third-party (relying party) validation of assertions related to resource allocations. By coupling the issuance of a certificate by a parent Certification Authority (CA) to the corresponding resource allocation, a test of the certificate's validity can be interpreted as validation of the associated resource allocation.

A Resource Certificate describes an action by the certificate issuer that binds a list of IP Address blocks and AS Numbers to the subject of the certificate. The binding is identified by the implicit association of the subject's private key with the subject's public key contained in the Resource Certificate, signed by the private key of the certificate's issuer. Any object signed by the subject's private key relates to an assertion of resource control, and can be validated via the matching public key contained in the certificate (and validation of the certificate itself in the context of the RPKI [14]).

The intent of the Resource Public Key Infrastructure (RPKI) is to support a hierarchy of X.509 certificates that allows relying parties to validate assertions about IP addresses and AS Numbers, and their use. The RPKI allows a relying party to determine if an address is valid to use in the context of the public Internet, and to validate assertions relating to the current "right-of-use" holder of an AS number or IP address.

The RPKI mirrors the resource allocation hierarchy. In this model the IANA issues certificates to each of the RIRs, describing in a resource extension to the certificate the complete set of number resources that have been allocated to that RIR. Each RIR issues certificates that correspond to allocations made by that RIR, where the resource extension to the certificate lists all the allocated resources, and the certificate holds the public key of the recipient of the resource allocation, signed with the private key of the RIR.

The common constraint within this PKI is that an issued certificate must contain a resource extension that contains a subset of the resources that are described in the resource extension of the issuing authority's certificate. This corresponds to the allocation constraint than an Internet Registry cannot allocate resources that were not allocated to the registry in the first place. The implication of this constraint is that if any party holds resources allocated from two or more registries then it will hold two or more resource certificates to describe the complete set of its resource holdings.

When an entity acquires an additional allocation, the associated certificate is reissued with a resource extension that matches the new allocation state. In the case of a reduction in allocated resources, the previously-issued certificate is revoked. In other cases there is no explicit revocation of the older certificates.

Validation of a certificate in the RPKI is similar to conventional certificate validation, establishing a chain of valid certificates, linked by issuer to subject, from a nominated trust anchor CA to the certificate in question. The additional constraint added by the RPKI is that every certificate in this validation path must be a valid resource certificate, and the resources described in the certificate are a subset of those described in the issuing authority's certificate.

The profile of Resource Certificates is described in Table 1, indicating all the fields that must be included in a resource certificate [15].

The distinguished name of the certificate's subject is normally nominated by the subject and verified by the issuer. In the RPKI the certificate issuer is not making any form of attestation regarding the right of the subject to assert any

TABLE I
RPKI CERTIFICATE PROFILE

Field	Value
Version	3
Serial Number	Positive integer, unique per issuer
Signature Algorithm	Minimum of SHA-256
Issuer	Distinguished Name of certificate issuer
Subject	Distinguished Name of Subject (Issuer-assigned)
Valid From / To	certificate validity dates
Subject Public Key Info	Subject's public key and algorithm
Basic Constraints	Intended use context
Subject Key identifier	SHA-1 hash of subject's public key
Authority Key Identifier	SHA-1 hash of issuer's public key
Key Usage	CA or EE certificate
CRL Distribution Point	URL of the CA's CRL
Authority Information Access	URL of the issuer's superior certificate
Subject Information Access	URL of the subject's repository publication point
Certificate Policies	Resource Certificate Policy Identifier
IP Resources	Issuer-allocated IPv4 and IPv6 addresses
AS Resources	Issuer-allocated AS numbers

particular identity. Consequently, in the RPKI the distinguished name is selected by the issuer, and is generated as a random string, so that it does not convey any particular identity of the subject, other than uniqueness within the name space used by the issuer.

All Resource Certificates must have the IP Addresses and/or AS Resources present, and marked as a critical extension. The contents of these extensions correspond exactly to the current state of IP address and AS number allocations from the issuer to the subject.

The current profiles for Resource Certificates and signed objects use a minimum of SHA-256 for the signature algorithm, and an RSA key size of 2048 bits. Experience has shown the wisdom in allowing for algorithm agility in such standard profiles, and while this choice of algorithm and key size represents a reasonable compromise between efficiency of use and cryptographic protection in the current environment, it is recognized that this pragmatic judgment will inevitably change over time, and stronger cryptographic algorithms and potentially longer key sizes will be required in the profile in the future. This consideration for algorithm agility has been incorporated in the RPKI profile.

Any holder of a resource who is in a position to make further allocations of resources to other parties must be in a position to issue Resource Certificates that correspond to these allocations. Similarly, any holder who wishes to use the RPKI to digitally sign an attestation needs to be able to issue an End Entity certificate to enable relying parties to validate such signatures. For this reason all issued certificates that correspond to resource allocations are CA certificates. Each CA certificate is capable of issuing subordinate CA certificates that correspond to further sub-allocations, and to issue (subordinate) EE certificates that enable verification of digital signatures on objects.

EE resource certificates are used in the RPKI to verify "with resources." For example, a resource holder may wish to authorize an AS to generate a route announcement for a particular address prefix. In this case the prefix holder would generate an EE resource certificate with the resource extension spanning the set of addresses that match the address prefixes that are the intended subject of the routing authority. It would place validity dates in the EE certificate that correspond to the intended validity dates of the routing authority. The authority object would contain the AS that is being authorized in this manner, and a description of the range of prefixes that the prefix holder has authorized, and the EE certificate. The object would be signed by the EE certificate's private key. A relying party could validate the authority to route by checking that the digital signature is correct, that the resources in the EE certificate encompass he prefixes specified in the document, and that the EE certificate itself is valid in the RPKI context.

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 must issue 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 revokes the corresponding EE certificate. It is also a property of this PKI that the key used to sign a CRL must be the same key used to sign the certificates being revoked, therefore binding a logical instance of a CA to a single key. Key rollover for a CA is performed by creating a new logical instance of the CA.

All Resource Certificates, CRLs, and other signed objects in the RPKI are published in openly accessible repositories. The set of all such repositories forms 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.

IV. SIGNED OBJECTS IN THE RPKI

The utility of a PKI lies in the ability to validate digitally signed information. The particular utility of the RPKI is not as means of validation of attestations of identity or role, but a means of validating the authority to use IP resources. While it is possible to digitally sign any digital object, it is proposed that the RPKI system uses a number of standard signed objects that have particular meaning in the context of routing security.

The common approach for all signed objects in the RPKI is to use a dedicated EE certificate to verify each object. In this way the issuer of the object can control the object's validity by having the ability to revoke the EE certificate at any time, so there is no need to create additional mechanisms within each signed object to control its validity. (The validity interval of the EE certificate also defines, implicitly, the lifetime of the signed objects it is used to verify.)

The first of these objects is the Route Origination

Authorization (ROA) [16]. A ROA is an authority, created by a prefix holder, than authorizes an AS to originate one or more specific route advertisements into the inter-domain routing system. A ROA is a digital object formatted according to the Cryptographic Message Syntax specification (CMS) [17] that contains a list of address prefixes and one AS number. The AS is the specific AS being authorized to originate a route advertisement, and the list of address prefixes are those that the AS is being authorized to originate. 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 match the IP address prefixes listed in the ROA's contents. As previously noted, the requirement in RPKI certificate issuance and validation is that Internet resources exactly follow allocation and assignment, and are a strict hierarchy. Therefore any valid subset of an RPKI 'branch' in the tree can be used to construct, and enable verification of, an exactly matching subset of address resources. EE certificate validation verifies rights to manage the resources, and requiring the resources to match the prefixes in the CMS associates these resources exactly with the ROA.

The ROA conveys a simple authority, and does not convey any further routing policy information, nor whether or not the AS holder has consented to actually undertake the routing action. The 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 digital signature, in a secure fashion.

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 (i.e., when only some legitimate route originations are authorized in a ROA), the absence of a ROA cannot be interpreted simply as invalid use of an address prefix. Similarly the presence of an invalid ROA should not necessarily invalidate a route advertisement in such a partial deployment scenario. (As an attacker could deliberately generate an invalid ROA for a route object that is otherwise valid, but not described in a valid ROA, it would be inappropriate for a BGP router to reject a route under such circumstances.)

This brief analysis shows why, during a partial deployment scenario, BGP routers need a 3-state model for route advertisements. Some routes will be valid (ROA-verified), some will be invalid, and others will be of unknown status.

If a given route matches exactly the information contained in a ROA whose EE certificate can be validated in the RPKI (a "valid" ROA) then the route can be regarded as a valid origination, and all other routes can be regarded as invalid. Where the prefix in a route is not described in any ROA and is not a more specific prefix of any ROA, then in a full deployment environment, such a route can be regarded as "invalid". In an environment of partial deployment, then a route that does not match any valid ROAs has an "unknown" validation outcome.

One way of feeding this information back into BGP is via a BGP LocalPref setting, where validated outcomes are more preferred, "unknown" validation credentials are essentially 'neutral', and "invalid" outcomes are less preferred, or can be rejected outright, depending on the local routing policy framework [18]. Care should be taken with such local policy settings relating to route rejection, as there is the consideration of potential circularity between the location of the repository containing the security credentials for a route object and the route object itself. If the repository publication point is located at an address protected by signed objects in that repository, then a relying party may need to accept an "invalid" route temporarily in order to access the security credentials that will validate the route.

While ROAs can be used to validate origination information, a related routing security question concerns the validity of the AS path information, that is, the sequence of AS's that describe the path from the origin to the recipient BGP router.

In attempting to validate an AS path there are a number of potential validation questions. The first and weakest question is: are all AS's in the AS Path valid AS's? A slightly stronger validation question is: do all the AS pairs in the AS Path represent valid AS adjacencies (where both AS's in the pair-wise association are willing to attest to the mutual adjacency). This latter validation question is used in the soBGP model of AS Path validation [8]. A yet stronger question is: do the sequence of AS's in the AS Path represent the actual propagation path of the BGP route object? This question is used as the basis of AS Path validation in the S-BGP model [7]. These differences of degree of path validation expose differences of approaches to AS path validation, and also expose to some extent the current uncertainty of the costs of path validation in operational environments. They also raise the question of what degree of validation outcomes can be achieved on a per-BGP Update processing level in BGP routers, and what can be validated externally to BGP and converted to simpler forms of update filters that are loaded onto routers. This is expected to remain an area of focus in the study of routing security for some time yet.

In looking at the AS adjacency question is it possible to construct an object similar in syntax to a ROA, that for a given AS lists all the adjacent AS's? One possible approach is through the use of AS Adjacency attestation Objects (AAO's) [19]. An AAO is a digitally signed object that provides a means of verifying an AS's attestation that it has a inter-domain routing adjacency with one or more AS's. In this instance, the RPKI validation relates to the holder of the attesting AS, so that an AAO is verified using an EE certificate issued under the signing AS (rather than any certificate associated with the list of AS's declared to be adjacent to this AS).

It would be reasonable for a relying party to infer from a single valid AAO that the signing AS may have the intent to advertise route objects across this adjacency, or may be prepared to learn route objects that are passed to it from the adjacent AS, or possibly both. However, an AAO is an asymmetric assertion, where one AS is claiming that an inter-

domain routing adjacency with another AS exists, but this claim is not explicitly acknowledged by the remote AS in the context of a single AAO. Relying parties may elect to place greater levels of confidence in the existence of an inter-domain routing adjacency when both AS's have signed and published AAO objects that contain mutual references. Like a ROA, an AAO is constructed using CMS as an envelope that carries an EE certificate and an ASN.1 description of the relevant AS numbers.

It is also possible to apply RPKI digital signatures to a set of IRR objects, using the principles of the RPSS [6] to guide the decision as to which party should sign the object [20]. RPSL "Aut-num" objects should be signed by the holder of the AS number, and RPSL "Inet-num" object should be signed by the holder of the IP address prefix. Under this model, an RPSL "Route" object should require the signature of both the AS holder and the IP address holder, signifying both the granting of an authority by the IP address holder, and the acceptance of this by the AS holder.

The advantage of using RPKI digital signatures in the context of an IRR is that it is then possible to divorce an IRR object from its point of publication, and allow relying parties to validate assertions relating to origination and routing policy with the strong assurance that the IRR objects are authentic and have not been altered in any way.

This approach would directly address the current weakness of the IRR dependency on the provenance of publication of IRR objects. Instead of weak trust in a "source" of IRR objects, a strong, and testable trust in the signatures would provide far greater assurance for relying parties that the IRR objects accurately represent the intentions and permissions of the object's maintainer.

V. Operating the RPKI

Almost all RPKI CA certificates, and all EE certificates, are regarded as relatively short-lived artifacts, i.e., regular re-issuance is normal and expected. Most PKIs that focus on identity rely on relatively long-lived certificates, and thereby can be designed to minimize overhead. Given the highly dynamic nature of routing (where it is not uncommon for several significant updates per day to be made to an ISP's routing model, either locally or globally) and its criticality to the stability of the Internet, the decision was made to not create very long-lived certificates. Instead, the RPKI model requires active management of current state, and frequent re-issuance of the EE certificates associated with signed objects. This, in turn, places an onus of responsibility on relying parties to perform regular sweeps across the distributed RPKI repository structure to ensure that the relying party is equipped with a local cache of up to date RPKI data. The question then arises as to the extent to which this model imposes a burden on both CAs and relying parties.

The number of participating entities in the RPKI is relatively modest in absolute terms, but when a daily (or more frequent) refresh cycle is taken into account the overall activity level could be significant It is expected that over time, a significant number of participants in Internet address management, which encompasses a population of the order 20,000 to 30,000 entities worldwide, will routinely publish RPKI CA certificates, (daily) CRLs, and signed objects (e.g., ROAs) for secure routing. This implies a potential population of discrete repository publication points of a similar order of magnitude.

The number of relying parties is also expected to be of the order of 10,000 entities. This corresponds, in very approximate terms, to the number of entities that provide BGP services as a transit AS, as distinct from stub AS domains. These relying parties would be expected to operate (loosely) synchronized local caches of the RPKI in order to perform validation checks on ROAs, and other signed objects to confirm the validity of routing information propagated through BGP.

To confirm that no changes have been undertaken at a publication point in the repository system the appropriate test is to see if the "manifest" for that point has changed. (A manifest is a digitally signed object that enumerates the names of all files at a publication point, and associates a hash value with each file. Every publication point in the repository must have an associated manifest.) If every relying party checks for changes every 24 hours, then each repository publication point would have its manifest polled by about 10,000 relying parties each day, and each relying party would need to check up to 30,000 manifests each day to see if the manifest, and hence the local repository data set, has changed. (If the manifest indicates any changes, the relying party will then need to fetch the objects that have changed.)

This synchronization load can be mitigated by the use of intermediaries that aggregate RPKI data into a single data collection and allow relying parties to synchronize against this single aggregate. It is anticipated that each RIR will maintain a repository that will consolidate RPKI repository data on behalf of its members, which would significantly reduce the number of distinct sites that need to be checked by a relying party.

Of course these estimates assume a stable steady state. In the even of some more fundamental change, perhaps as a result of some forced re-keying across a large proportion of the RPKI, then the synchronization load would also need to include the downloading of all the certificates and signed objects by each relying party. This would represent a significantly higher load than the simple repository freshness check.

The other operational aspect is a design decision as to whether the validation of a route object is performed by a router using the router's processing capability, or whether cryptographic processing can be off-loaded to a dedicated system. For origin validation, the later approach is currently anticipated. Specifically, it is anticipated that each ISP will operate a server that fetches RPKI repository data, processes it, and makes it available to the BGP routers in the AS of the ISP. For AS path validation it is not clear that a similar, offline validation approach is appropriate. The concern is that while the binding of an origin AS to a prefix typically changes infrequently, whereas AS paths may change very quickly. The next section discusses this issue in more detail.

VI. Current Status and Next Steps for the RPKI

Resource Certificates and the associated RPKI represent a major component of the effort to construct a secure inter-domain routing framework. The use of a clear and explicit structure to validate attestations regarding the control of address resources and their use in the context of routing allows for a simpler security framework and more efficient problem resolution than existing approaches. An RPKI, even partially populated with signed information, allows BGP speakers to make preferential selections to use routing information where the IP address block and the AS numbers being used are recognized as being valid to use, and that the parties using these IP addresses and AS numbers are properly authorized to so do. The RPKI can also efficiently identify instances of unauthorized use of IP addresses and attempts to hijack routes, and do so in a uniform manner, rather then via a more haphazard sequence of arbitrary local policy decisions being made, based on varying information and local assumptions.

However, the RPKI represents only one part of a larger framework of securing inter-domain routing, and the next step is that of applying the RPKI to the local BGP processing framework. Within the agenda of the current standardization effort in the Internet Engineering Task Force, there is also the need to move beyond validation of route origination and look at the associated issue of validation of the AS Path. This is a more challenging task of attempting to validate whether transitive information as presented in the AS Path represents the path used by the routing system. Path validation would provide a greater level of confidence that the initial forwarding hop associated with an offered route represents the correct first hop along a useable forwarding path for packets to reach the network destination.

The issues in path and forwarding validation include not only a consideration of what can be secured and validated, but also issues of scalability and efficiency in terms of deployment cost. The various approaches to path security studied so far vary widely in terms of the amount of routing information that is validated, the level of trust that can be placed in a validation outcome, and the storage and computational overhead of generating and validating digital signatures on routing information. As an example of the tradeoffs involved consider the choice of signature algorithms. The RPKI has elected to use RSA as the signature algorithm for certificates, CRL, and other signed objects. This choice is reasonable be cause these objects are signed offline and likely can be validated outside of a router, even for path validation. However, the S-BGP approach to path validation proposed using DSA for some signed objects, to reduce the size of the signatures that would be passed in BGP Update messages. The use of different signature algorithms for different classes of objects can be accommodate by the RPKI. The next step in securing BGP probably will entail exploring path security options relative to the storage and computational overhead imposed on routers, and the security offered by each option.

VII. Conclusion

Novel features incorporated into the PKI include the avoidance of any attestation relating to the identity or role of the subject in a Resource Certificate, and re-purpose the certificate to a "right-of-use" for IP number resources. The distributed repository system is protected through the use of the manifest construct, to allow relying parties to detect if their access to a repository has been corrupted in any way. The validity of signed objects is controlled by the use of a dedicated EE certificate for each signed object, allowing an authority to be revoked through the conventional use of a CRL.

The RPKI has been designed as a robust, 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.

References

[1] Recommendation X.509: The Directory Authentication Framework, ITU-T, 2000.

[2] L. Daigle, "Whois Protocol Specification," *Request For Comment* RFC3912, September 2004.

[3] C. Alaettinoglu, et al., "Routing Policy Specification Language(RPSL)," *Request for Comment* RFC2622, June 1999.

[4] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," *Request for Comment* RFC271, January 2006.

[5] R. Steenbergen, "Examining the Validity of IRR Data," *NANOG 44*, October 2008.

[6] C. Villamizar, et al., " Routing Policy System Security," *Request For Comment* RFC2725, December 1999.

[7] 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.

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

[9] 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.

[10] 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'03)*, February 2003.

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

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

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

[14] M. Lepinski, S. Kent, "An Infrastructure to Support Secure Internet Routing," work in progress, February 2008.

[15] G. Huston, G. Michaelson, R. Loomans, "A Profile for X.509 PKIX Resource Certificates," work in progress, September 2008.

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

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

[18] G. Huston, G. Michaelson, "Validation of Route Origination in BGP using the Resource Certificate PKI," work in progress, August 2008.

[19] G. Huston, G. Michaelson, " A Profile for AS Adjacency Attestation Objects," work in progress, May 2009.

[20] R. Kisteleki, J. Boumans, "Securing RPSL Objects with RPKI Signatures," work in progress, October 2008.